

Pipeline 130 User's Guide

Ascend Communications, Inc.

Part Number: 7820-0278-003

For software version 6.0

January 30, 1998

Pipeline, MAX, and Bandwidth-on-Demand are trademarks, and Ascend and the Ascend logo are registered trademarks of Ascend Communications, Inc. Other trademarks and trade names mentioned in this publication belong to their respective owners.

Copyright © 1998, 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 Customer Service

You can request technical assistance or additional information by telephone, email, fax, or modem, or over the Internet.

Obtaining Technical Assistance

If you need technical assistance, first gather the information that Ascend Customer Service will need for diagnosing your problem. Then select the most convenient method of contacting Ascend Customer Service.

Information you will need

Before contacting Ascend Customer Service, gather the following information:

- Product name and model
- Software and hardware options
- Software version
- Service Profile Identifiers (SPIDs) associated with your product
- Your local telephone company's switch type and operating mode, such as AT&T 5ESS Custom or Northern Telecom National ISDN-1
- Whether you are routing or bridging
- Type of computer you are using
- Description of the problem

How to contact Ascend Customer Service

After you gather the necessary information, contact Ascend in one of the following ways:

Telephone in the United States	800-ASCEND-4 (800-272-3634)
Telephone outside the United States	510-769-8027 (800-697-4772)
Austria/Germany/Switzerland	(+33) 492 96 5672
Benelux	(+33) 492 96 5674

France	(+33) 492 96 5673
Italy	(+33) 492 96 5676
Japan	(+81) 3 5325 7397
Middle East/Africa	(+33) 492 96 5679
Scandinavia	(+33) 492 96 5677
Spain/Portugal	(+33) 492 96 5675
UK	(+33) 492 96 5671
Email	support@ascend.com
Email (outside US)	EMEAsupport@ascend.com
Facsimile (FAX)	510-814-2312
Customer Support BBS by modem	510-814-2302

You can also contact the Ascend main office by dialing 510-769-6001, or you can write to Ascend at the following address:

Ascend Communications, Inc.,
One Ascend Plaza,
1701 Harbor Bay Parkway,
Alameda, CA 94502

Need information about new features and products?

Ascend is committed to constant product improvement. You can find out about new features and other improvements as follows:

- For the latest information about the Ascend product line, visit our site on the World Wide Web:

<http://www.ascend.com>

- For software upgrades, release notes, and addenda to this manual, visit our FTP site:

<ftp.ascend.com>

Contents

Ascend Customer Service	iii
About This Guide	xix
How to use this guide	xix
What you should know	xix
Documentation conventions	xx
How to use the on-board software	xxi
Changing parameter values	xxii
Saving or discarding your changes	xxii
Manual set	xxii
Configuring WAN Connections	1-1
About Wide Area Network (WAN) connections	1-1
Link encapsulation	1-2
Nailed groups	1-3
How calls are initiated	1-4
How calls are answered	1-5
Data compression options	1-6
The Answer profile	1-7
Connection profiles	1-9
Session options	1-11
Telco options	1-13
Encapsulation options	1-14
PPP connections	1-14
MP, MPP, and MP+ connections	1-16
Dynamic bandwidth allocation (DBA)	1-17
Guidelines for configuring DBA	1-18

Monitoring DBA	1-19
Clearing a call on the basis of idle bandwidth	1-19
BACP connections	1-22
Nailed MPP connections	1-23
Configuring Frame Relay connections	1-25
Configuring a Frame Relay profile	1-28
Configuring a gateway connection	1-30
Inverse ARP for Frame Relay	1-31
Configuring a T1 line	1-32
Configuration steps	1-32
Backup Connection disconnect timer	1-35
Configuring a V.35 serial WAN port	1-35

Configuring IP Routing 2-1

Introduction to IP routing on the Pipeline	2-1
Host-to-router connections	2-2
Router-to-router connections	2-3
Subnet mask notation	2-4
IP routing in the Answer profile	2-7
Connection profiles and IP routes	2-8
How the Pipeline uses its routing table	2-8
RIP-v2 and RIP-v1 routing	2-9
Interface-based routing	2-10
System behavior with a numbered interface	2-11
Configuring interface-based routing	2-11
Specifying the remote interface address	2-12
Multicast forwarding and IGMP functionality	2-13
Managing the routing table	2-14
Parameters that affect the routing table	2-14
Static and dynamic routes	2-16
Configuring static routes	2-16
Creating a Static Rtes profile	2-18
Configuring the default route	2-19
Specifying default routes on a per-user basis	2-20
Enabling the Pipeline to use dynamic routing	2-21
If you are using RIP-v1	2-21
Configuring RIP-v2 on Ethernet	2-22
Configuring RIP for incoming WAN connections	2-22
Configuring RIP for a particular connection	2-23

Route preferences	2-24
Viewing the routing table	2-25
Fields in the routing table	2-27
Removing down routes to a host	2-29
Identifying Temporary routes in the routing table	2-30
Configuring IP routing connections	2-30
Checking remote host requirements	2-30
Example host connection with static address	2-31
Example router connection	2-33
Example router connection on a subnet	2-35
Ascend Tunnel Management Protocol (ATMP)	2-37
Using a Pipeline in a virtual private network	2-37
Foreign and home agents	2-37
Configuring a home agent in router mode	2-38

IP Address Management 3-1

Connecting to a local IP network	3-1
Assigning the Ethernet interface IP address	3-3
Creating a subnet for the Pipeline	3-3
Assigning two addresses: Dual IP	3-4
Using Ping to verify the address	3-6
Enabling proxy mode in the Pipeline	3-6
Enabling DNS on the Pipeline	3-7
Generating UDP checksums	3-7
Updating other routers on the backbone	3-8
BOOTP Relay	3-9
DHCP services	3-10
How IP addresses are assigned	3-10
Configuring DHCP services	3-11
Setting up a DHCP server	3-14
Setting up Plug and Play support	3-14
Setting up DHCP spoofing	3-14
Dial-in user DNS server assignments	3-15
Configuring DNS servers in the Ethernet profile	3-15
Configuring DNS servers in the Connection profile	3-16
Local DNS host address table	3-17
Configuring the local DNS table	3-18
Creating the local DNS table	3-19
Editing the local DNS table	3-20

Deleting an entry from the local DNS table	3-21
Restrictions for names in the local DNS table	3-21
User-definable TCP connection retry timeout	3-21
Network Address Translation (NAT) for a LAN	3-22
Single-address NAT and port routing	3-23
Outgoing connection address translation	3-23
Incoming connection address translation	3-24
Translation table size	3-24
Multiple-address NAT	3-25
Configuring single or multiple address NAT	3-26
NAT for Frame Relay	3-28
Configuring NAT port routing (Static Mapping submenu)	3-29
Routing all incoming sessions to the default server	3-29
Routing incoming sessions for up to 10 servers on a LAN	3-30
Disabling routing for specific ports	3-32
Well-known ports	3-32

Configuring IPX Routing 4-1

How the Pipeline performs IPX routing	4-1
IPX Service Advertising Protocol (SAP) tables	4-2
IPX Routing Information Protocol (RIP) tables	4-3
Extensions to standard IPX	4-3
Virtual IPX network for dial-in clients	4-4
Optimized access for dial-in NetWare clients	4-4
IPX Route profiles	4-6
IPX SAP filters	4-7
Dial Query	4-8
Watchdog spoofing	4-9
Automatic SPX spoofing	4-9
WAN considerations for NetWare client software	4-10
IPX in the Answer profile	4-11
Adding the Pipeline to the local IPX network	4-12
Checking local NetWare configurations	4-12
Configuring IPX on the Pipeline Ethernet interface	4-13
Using IPXping to check the configuration	4-14
Defining a virtual IPX network for dial-in clients	4-15
Working with the RIP and SAP tables	4-15
Viewing the RIP and SAP tables	4-16
Configuring RIP in a Connection profile	4-17

Configuring a static IPX route	4-18
Configuring SAP in a Connection profile	4-20
Managing IPX SAP filters	4-21
Defining an IPX SAP filter	4-21
Applying an IPX SAP filter	4-23
Configuring IPX routing connections	4-24
An example dial-in client connection	4-24
An example with NetWare servers on both sides of the link	4-25
An example with local NetWare servers only	4-29

Configuring the Pipeline as a Bridge 5-1

Introduction to Ascend bridging	5-1
How a bridged WAN connection is initiated	5-2
Physical addresses and the bridge table	5-2
Broadcast addresses and Dial Brdcast	5-3
How bridged connections are established	5-3
Bridging in the Answer profile	5-4
About IPX bridging	5-4
When there is no server support on the local network	5-5
When there is no server support on the remote network	5-5
When there is server support on both networks	5-5
IPX routing and bridging on the same connection	5-5
Enabling bridging	5-6
Managing the bridge table	5-7
Parameters that affect the bridge table	5-7
Transparent bridging	5-7
Static bridge-table entries	5-8
Configuring bridged connections	5-9
An example AppleTalk bridged connection	5-10
An example IPX client bridge (local clients)	5-13
An example IPX server bridge (local servers)	5-14
An example IP bridged connection	5-16

Defining Filters and Firewalls 6-1

Introduction to filters	6-1
Data filters for dropping or forwarding certain packets	6-2
Call filters for managing connections	6-4
Predefined call filters	6-5

Overview of Filter profiles	6-6
Filtering inbound and outbound packets	6-7
Selecting filter type and activating the filter	6-8
Defining generic filter conditions	6-9
Defining IP filter conditions	6-10
Example filters	6-12
An example generic filter to handle AppleTalk broadcasts	6-12
An example IP filter to prevent address spoofing	6-16
An example IP filter for more complex security issues	6-19
Working with predefined call filters	6-21
NetWare Call filter	6-21
Extending the predefined filter for RIP packets	6-23
Defining a SNEP data filter for Ethernet	6-24
IP Call filter	6-26
AppleTalk Call filter	6-26
Display unwanted dial-out packets	6-28
When packets are not captured	6-28
Turning on the diagnostic option	6-29
Displaying packets	6-29
Secure Access Firewalls	6-34
Determining if Secure Access is present	6-34
Firewall profiles	6-34
Assigning firewalls to a Connection profile	6-35
Assigning firewalls to the Mod Config profile	6-35
Filter persistence	6-36
Background on firewall and filter persistence	6-36
Filter persistence and Connection profiles	6-37

Setting Up Pipeline Security 7-1

Recommended security measures	7-1
Changing the Full Access security level password	7-3
Activating the Full Access security level	7-4
Making the Default security level restrictive	7-4
Assigning a Telnet password	7-5
Changing the SNMP read and write community string	7-5
Requiring profiles for incoming connections	7-7
Turning off ICMP redirects	7-7
Pipeline Security profiles	7-7
Default security level	7-8

Security profile passwords	7-8
Security privileges	7-8
Using the Full Access profile	7-9
Defining a second Security profile	7-10
Connection security	7-11
Authentication protocols	7-12
Name and password verification	7-12
Calling-line ID authentication	7-13
Settable disconnect cause codes for CLID authentication	7-14
Callback security	7-15
Expect callback support	7-16
Using filters to secure the network	7-16
Using security cards	7-17
Supporting outbound security card calls	7-18
Configuring the Pipeline to recognize the APP Server utility	7-21
Invoking password mode in the Pipeline	7-22

Pipeline System Administration 8-1

Overview of administration functions	8-1
Activating administrative privileges	8-3
Configuring administration options	8-4
Setting system values	8-4
Configuring the Pipeline to interact with syslog	8-5
Syslog messages	8-7
Using the Pipeline status windows	8-9
Performing system administration operations	8-10
Using DO commands	8-10
Saving the Pipeline configuration	8-11
Restoring the Pipeline configuration	8-14
Resetting the Pipeline	8-15
Using the terminal server interface	8-17
Invoking and quitting the terminal server interface	8-17
Terminal server commands	8-17
Accessing a local Pipeline via Telnet	8-21

APP Server utility A-1

About the APP Server utility	A-1
APP Server installation and setup	A-2

Configuring the Pipeline to use the APP server	A-2
Using App Server with Axent SecureNet	A-3
Creating banner text for the password prompt	A-3
Installing and using the UNIX APP Server	A-6
Installing and using the APP Server utility for DOS	A-8
Installing and using the APP Server utility for Windows	A-10
Installing the APP Server utility for Windows 3.1	A-11
Installing the APP Server utility for Windows 95	A-12
Installing the APP Server utility for Windows NT	A-12
Installing APP Server on a Macintosh	A-13
 Troubleshooting	 B-1
Cabling problems: Rule these out first	B-1
Common problems and their solutions	B-2
General problems	B-2
Profile configuration problems	B-2
Hardware configuration problems	B-3
Problems configuring the Pipeline	B-5
ISDN BRI interface problems	B-8
T1 and ISDN BRI circuit-quality problems	B-10
T1 access problems	B-11
Bridge/router problems	B-12
Problems accessing the remote network	B-13
Check the installation	B-13
Configuration problems	B-14
Traps for BRI linkUp and linkDown	B-16
 Upgrading system software	 C-1
What you need to upgrade system software	C-1
Displaying the software load name	C-2
The upgrade procedure	C-3
Activating a Security Profile	C-4
Guidelines for upgrading system software	C-5
Before you begin	C-6
Upgrading system software with a standard load	C-7
Upgrading using the serial console	C-7
Upgrading standard load using TFTP	C-8
Upgrading system software to a fat or extended load	C-9

Recovering from a failed upgrade	C-10
Pipeline checks compatibility of downloaded files	C-11

Glossary	Glossary-1
-----------------------	-------------------

Index	Index-1
--------------------	----------------

Figures

Figure 1-1 Bandwidth algorithms for MP+ calls.....	1-18
Figure 1-2 Gateway connections to the Frame Relay network	1-27
Figure 2-1 An IP routing connection between two networks.....	2-3
Figure 2-2 A class C address	2-4
Figure 2-3 A 29-bit netmask and number of supported hosts	2-5
Figure 2-4 An IP routing connection serving as a static route	2-17
Figure 2-5 When a two-hop static route is required with RIP off.....	2-17
Figure 2-6 A dial-in user requiring a static IP address (a host route).....	2-32
Figure 2-7 A router-to-router IP connection	2-33
Figure 2-8 A connection between local and remote subnets.....	2-35
Figure 3-1 Creating a subnet for the Pipeline.....	3-3
Figure 3-2 Dual IP and shared subnet routing.....	3-5
Figure 4-1 A dial-in client requiring dynamic IPX network assignment...	4-24
Figure 4-2 A connection with NetWare servers on both sides	4-25
Figure 4-3 A dial-in client that belongs to its own IPX network	4-29
Figure 5-1 Negotiating a bridge connection (PPP encapsulation).....	5-3
Figure 5-2 How the Pipeline creates a bridging table	5-8
Figure 5-3 An example IPX client bridging connection	5-13
Figure 5-4 An example IPX server bridging connection.....	5-14
Figure 5-5 lAn example IP bridging connection	5-16
Figure 6-1 Data filters can drop or forward certain packets.....	6-2
Figure 6-2 Call filters used to prevent resetting the timer.....	6-4
Figure 6-3 Filter organization and terminology	6-6
Figure 7-1 RADIUS acting as client of ACE or Safeword server.....	7-18
Figure 8-1 Status windows	8-9

Tables

Table 1-1	Frame Relay and gateway profiles	1-27
Table 2-1	IP address classes and default netmasks	2-4
Table 2-2	Standard netmasks and Ascend netmask notation	2-5
Table 8-1	Terminal server commands.....	8-18
Table A-1	APP Server INI file contents.....	A-4
Table C-1	Format of binary loads (size comparisons)	C-5

About This Guide

How to use this guide

This manual is part of a set that describes all the standard features of a Pipeline running software version 6.0. Some features might not be available with older versions or specialty loads of the software. Features available only with specialty loads are documented in separate publications.

This manual is organized with basic information about setting up connections first, followed by more specific information about administering the unit.

Read this manual to find out how to create more connections to remote sites, or to refine the way traffic is handled by the Pipeline. If you only need to connect to a single network, and your primary connection profile is set up, use this manual to secure the Pipeline. See Chapter 7, “Setting Up Pipeline Security.”

If you are a network administrator, use this manual to troubleshoot connection problems, set up filters, set authentication methods, manage local or remote units, and upgrade your unit’s onboard software.

Refer to the *Reference Guide* for information about possible values for any setting, for examples, and to find out which settings depend on others when enabling features.

What you should know



To configure the Pipeline, you need to understand the following:

- Internet or telecommuting concepts

- Wide area network (WAN) concepts
- Local area network (LAN) concepts, if applicable

Documentation conventions

The following list explains how special characters and typographical conventions are used in this manual.

Convention	Meaning
Monospace text	Represents text that appears on your computer’s screen, or that could appear on your computer’s screen.
Boldface monospace text	Represents characters that you enter exactly as shown (unless the characters are also in <i>italics</i> —see <i>Italics</i> , below). If you could enter the characters, but are not specifically instructed to, they do not appear in boldface.
[]	Square brackets indicate an optional argument you might add to a command. To include such an argument, type only the information inside the brackets. Do not type the brackets unless they appear in bold type.
>	Points to the next level in the path to a parameter. The parameter that follows the angle bracket is one of the options that appears when you select the parameter that precedes the angle bracket.
<i>italics</i>	Italics represent variable information. Do not enter the words themselves; enter the information they represent.
Press Enter	Means press the Enter, or Return, key or its equivalent on your computer.
Note:	Introduces important additional information.
 Caution:	Warns that a failure to follow the recommended procedure could result in loss of data or damage to equipment.
 Warning:	Warns that a failure to take appropriate safety precautions could result in physical injury.

How to use the on-board software

This manual describes how to change the settings in the on-board software to add, change, or remove functions on the Pipeline. You can access the on-board software in these ways:

- By Telneting to the unit using the IP address to make the connection.
- By first establishing a serial connection from the terminal port on the back of the Pipeline to a COM port on your computer, and then by using VT100 terminal emulation software to access the on-board software.

The on-board software looks similar to this:

Main Edit Menu >Configure... 00-000 System 00-000 Ethernet 30-000 Serial WAN	10-100 1 Link D B1 - B2 -	00-200 17:20:50 > Line Ch
	20-100 Sessions > 0 Active	20-500 DYN Stat Qual N/A 00:00:00 OK 0 channels CLU 0% ALU 0%
	20-300 WAN Stat >Rx Pkt: Tx Pkt: CRC:	20-400 Ether Stat >Rx Pkt: Tx Pkt: Col:
	00-100 Sys Option >Security Prof: Software S/N:	00-400 HW Config >BRI Interface Adrs: Enet I/F: UTP

The Main Edit Menu (the window at the far left) is where you add, change, or remove settings. The other windows (in the middle and far right columns) are the status windows. Some status windows contain lists of information. Use the tab key to move from window to window, and use the up and down arrow keys or Ctrl-N (next) or Ctrl-P (previous) to scroll through the lists and menus. To open a menu, place the cursor (>) next to the menu name and press Enter.

With the exception of parameters designated N/A (not applicable), you can edit all parameters in any menu. N/A means the parameter is dependent on another parameter that is set to a value that is causing this parameter to not be used. (See the *Reference Guide* for dependency information.)

Changing parameter values

When a parameter has preset choices, press Enter to cycle through the choices. To select the current value, use the arrow key or Ctrl-N to move to the next field.

To edit a text-based parameter, move the cursor to the parameter and press Enter. An edit field opens, delimited by brackets, such as these []. A blinking text cursor appears in the brackets, indicating that you can start typing text. If the field already contains text, it is cleared when you type a character. To modify only a few characters of existing text, use the arrow keys to position the cursor, then delete or overwrite the characters. To close the edit field and accept the entry, press Enter.

Saving or discarding your changes

To save your changes and exit the menus, press the Esc key. If you have changed any parameter, the Exit menu appears, and provides choices to accept your changes and exit, or discard your changes and exit.

Manual set

This manual is part of a set that includes the following publications:

- *Pipeline Start Here* explains how to install the *Pipeline*, how to use the Java-based Pipeline Configurator, how to use the on-board software, and how to set up your primary connection.
- *Pipeline User's Guide* explains how to configure the *Pipeline* as a router or bridge, and how to manage the inbound and outbound traffic over the unit.
- *Pipeline Reference Guide* contains alphabetical listings of all the parameters and all of the fields in the status menus, and a section that explains how to use the DO commands.

Configuring WAN Connections

This chapter contains the following topics:

About Wide Area Network (WAN) connections	1-1
Link encapsulation	1-2
Nailed groups	1-3
How calls are initiated	1-4
How calls are answered	1-5
Data compression options.	1-6
The Answer profile.	1-7
Connection profiles	1-9
Configuring Frame Relay connections.	1-25
Configuring a T1 line	1-32
Configuring a V.35 serial WAN port	1-35

About Wide Area Network (WAN) connections

In order to connect to the wide area network, the Pipeline needs to know what attributes to apply to each incoming or outgoing link. For example, it needs to know how to negotiate the initial handshake with the remote end, what kind of authentication is required, what kind of compression needs to be agreed upon,

what data rates are available end-to-end, how much bandwidth can be allocated and which end will add it, what kind of encapsulation can be supported, and other information.

This chapter explains how to set up the Answer and Connection profiles. A profile is a group of settings that define the attributes needed to set up or answer a call. You can define multiple Connection profiles and one Answer profile. Connection profiles are used for both incoming and outgoing connections. The Answer profile supplies general setup information that can be used to reject or set up a connection if there is no Connection profile matching the caller's settings.

The first Connection profile is automatically created when you set the parameters in the Configure menu, which is described in the *Start Here* booklet.

Link encapsulation

One of the main agreements between the caller and the answering device must be the type of link encapsulation used. The caller must encapsulate all outbound packets before sending them across the WAN, and the answering device must unencapsulate them before forwarding the packets to the local network. Following are the types of link encapsulation supported by the Pipeline:

Method	Connection description and attributes
PPP	<p>Point-to-Point Protocol (PPP), is a single-channel connection that connects to any other device running PPP.</p> <p>PPP connections support password authentication using Password Authentication Protocol (PAP), Challenge Handshake Authentication Protocol (CHAP), and Microsoft CHAP (MS-CHAP).</p> <p>They can support IP routing, IPX routing, or protocol-independent bridged connections.</p> <p>They can be dial-in or dial-out switched connections.</p>
MP	<p>Multilink PPP (MP) call using two channels. This type of connection uses PPP to initiate the call and to add a second channel. Once the second channel is connected, it is used for the duration of the call. It is not as flexible as MPP or MP+.</p>

Method	Connection description and attributes
MPP (includes MP+)	<p>Multichannel Point-to-Point Protocol (MPP), Multilink Protocol Plus (MP+), and Bandwidth Allocation Control Protocol (BACP) are all enhancements to PPP for supporting multi-channel links. (MP+ is an implementation of MPP developed by Ascend Communications, Inc.)</p> <p>If a connection is set up for MPP, the Pipeline first requests MP+. If the other side of the connection doesn't support MP+, the Pipeline requests MP. If that protocol is refused, PPP is used instead. That is why the term "PPP connection" is often used to mean any one of these encapsulation methods when the number of channels is not relevant.</p>
Frame Relay	<p>The Frame Relay RFC 1490 standard does not support authentication.</p> <p>A Frame Relay gateway connection supports routing and bridging to and from the switch across a nailed connection.</p> <p>Some Ascend units provide Frame Relay operations as a software option.</p>

Nailed groups

A nailed connection is a permanent, physical circuit that is always up as long as the physical connection persists. (A nailed connection can also be a permanent virtual circuit, which is not a single physical connection, but a dedicated, switched link.) If the unit or central switch resets or if the link is terminated for any reason, the Pipeline attempts to restore the link at 10-second intervals. If the Pipeline or the far-end unit is powered off, the link is restored when power is restored.

On an ISDN line, a nailed connection uses one of the line's channels. A nailed T1 line or a Frame Relay link is not channelized and is always 100% nailed up.

To make channels available for a nailed connection, you have to designate them for nailed usage by assigning them to a group number.

Note: Make sure the group numbers are unique across all WAN interfaces.

Configuring WAN Connections

How calls are initiated

The group numbers for the Pipeline WAN interfaces are as follows:

- If you set the Chan Usage parameter to Leased/Switch, the Group number for the first B channel is 1 (this value cannot be changed).
- If you set the Chan Usage parameter to Switch/Leased, the Group number for the second B channel is 2 (this value cannot be changed).
- For a Nailed T1 line, serial WAN, or DDS-56, set the group number to 3.

Assign group numbers to channels as follows:

- For PPP encapsulated connections to other routers or bridges, set the group number specifying a value for the Ethernet > Connection > *any profile* > Telco > Group parameter.
- For Frame Relay encapsulated connections, set the group number by specifying a value for the Ethernet > Frame Relay > *profile* > Nailed Grp parameter.
- For a Nailed T1 line, set the group number by specifying a value for the Nailed T1 > Mod Config > Nailed T1 Group parameter.

How calls are initiated

When configured to bridge, the Pipeline initiates a bridged connection across the Wide Area Network (that is, it calls out) whenever it receives a broadcast packet or a packet whose destination is not on the local LAN. When configured to route, the Pipeline initiates a connection when it has a route to the destination. But the Pipeline does not attempt to route every packet unless it is configured as the default gateway.

You can manually dial a connection from a connection profile by pressing Control-D to invoke the DO menu, then selecting dial. (For further discussion of manual dialing, see the *Reference Guide* chapter on using the DO commands.)

For more information on how the Pipeline initiates calls, see:

- Chapter 2, “Configuring IP Routing.”
- Chapter 4, “Configuring IPX Routing.”
- Chapter 5, “Configuring the Pipeline as a Bridge.”

You can control how the Pipeline brings up WAN sessions using these methods:

- Use filters to block certain packets, such as broadcast, or IPX RIP or SAP queries, from bringing up a connection to the remote network. (For information about creating filters, refer to Chapter 6, “Defining Filters and Firewalls.”)
- When bridging, you can prevent the Pipeline from dialing out when it receives broadcasts by setting the Dial Brdcast parameter to No.
- When routing IPX, you can prevent the Pipeline from dialing out when it receives IPX queries by setting the Dial Query parameter to No.

How calls are answered

Before the Pipeline answers an incoming call, it checks the Answer profile to see if Calling Line ID (Id Auth) authentication is required. Id Auth verifies the caller’s phone number before answering the call. If Id Auth authentication is required and the phone number doesn’t match a Connection profile, the Pipeline drops the call. If Id Auth is not required or if a matching Connection profile is found, the Pipeline answers the call and applies the following tests:

1 Is the encapsulation type available?

The Pipeline supports PPP, MP+ (MPP), MP, BACP, and Frame Relay. If a call does not use the encapsulation type specified in the Connection profile, or cannot use PPP or Frame Relay, the Pipeline drops the call.

2 Is authentication required?

For PPP or MP+ calls, the Answer profile’s Recv Auth parameter might require PAP, CHAP, or MS-CHAP.

Frame relay does not support call authentication.

If authentication is required, a matching Connection profile must be found. If Answer > Profile Reqd=No and Id Auth=Ignore, the Answer profile parameters are used to build the connection.

3 Is there a matching Connection profile?

The Pipeline can accept a call defined in a Connection profile if the Connection > *any profile* > Telco Options > AnsOrig parameter is set to Both or Ans Only. The default is Both.

The Pipeline attempts to match the caller’s name and password to a Connection profile. If password authentication is not required, the Pipeline

can match IP-routing PPP calls against the IP address specified in the LAN Adrs parameter of the Connection profile.

4 What information is used to build the connection?

If authentication succeeds, the Pipeline builds the connection with the encapsulation, Telco Options, and Session Options specified in the Connection profile. If you configure the Pipeline to ignore authentication and do not require a Connection profile, the Pipeline uses the Answer profile to build the connection.

When the connection is established, the Pipeline forwards the call to its bridge or router software and begins routing or bridging the packets.

Data compression options

For data compression to take effect, both sides of a connection must support it. The Pipeline supports the following types of data compression:

Compression	Description
Stac	<p>For PPP-encapsulated calls, refers to a pre-RFC implementation of the Stacker compression algorithm, developed by Stac Electronics, Inc., which modifies the standard LZS compression algorithm to optimize for speed (as opposed to optimizing for compression).</p> <p>Stac compression is one of the parameters negotiated when setting up a PPP connection.</p>
Stac-9	<p>Requests the standard Stac compression described by the Stac RFC. If you chose to use Stac compression, set Link Compression for MS-Stac or Stac-9. If the far-end of the link does not accept MS-Stac or Stac-9, your unit will try to set up compression corresponding to the Stac setting. If this compression also fails, the unit runs the link uncompressed.</p>

Compression	Description
MS-Stac	<p>For PPP-encapsulated calls. MS-Stac refers to Microsoft LZS Coherency compression for Windows 95. This is a proprietary compression scheme for Windows 95 (not Windows NT).</p> <p>If the caller requests MS-Stac and the matching profile does not specify MS-Stac compression, the connection appears to come up correctly but no data is routed. If the profile is configured with MS-Stac and the caller does not acknowledge that compression scheme, the Pipeline attempts to use standard Stac compression, and if that doesn't work, it uses no compression.</p>
VJ Comp	<p>For TCP/IP connections. VJ Comp applies only to packets in TCP applications, such as Telnet. When you turn it on, the Pipeline applies TCP/IP header compression for both ends of the link.</p>

The Answer profile

Answer profiles contain parameters to build connections for incoming callers. Before the Pipeline answers an incoming call, it checks the settings in its Ethernet > Answer profile for information about what to do. If the call does not include the information required by the Answer profile, the Pipeline hangs up.

If the call includes the required information, the Pipeline looks for a matching Connection profile. If it finds one, it uses information in the Connection profile to set up the call. If a match is not found, the Answer profile specifies how to build the connection.

Note: The parameter Ethernet > Answer > Profile Req'd must be set to No to build a connection for a call that does not have a matching Connection profile.

To set up a basic Answer profile:

- 1 Open the Ethernet > Answer profile. The following menu is an example:

```
Force 56=No
Profile Req'd=No
Id Auth=Ignore
```

Configuring WAN Connections

The Answer profile

PPP options...

Session options...

- 2 To require a matching profile for incoming calls, set Profile Req'd=Yes.

This prevents the Pipeline from building a connection on the basis of parameters in the Answer profile.

- 3 If appropriate, set Id Auth=Required.

Some connection types do not provide an authentication method. If you plan to allow those types of calls, you might need to use Id Auth (calling line ID). (For more information, see Chapter 7, "Setting Up Pipeline Security.")

Next, set options for PPP and MP+ calls:

- 1 Open the PPP Options submenu.

Route IP=Yes

Route IPX=Yes

Bridge=Yes

Recv Auth=Either

MRU=1524

LQM=No

LQM Min=600

LQM Max=600

Link Comp=Stac

VJ Comp=Yes

Dyn Alg=Quadratic

Sec History=15

Add Pers=5

Sub Pers=10

Min Ch Count=1

Max Ch Count=1

Target Util=70

Idle Pct=0

- 2 Turn on routing and bridging for the connection as appropriate.

For example:

```
Route IP=Yes
Route IPX=No
Bridge=No
```

Note: You must have routing or bridging globally enabled in the Ethernet > Mod Config menu or in the Configure menu in order to route or bridge in a Connection profile.

- 3 Set the Recv Auth parameter to PAP, CHAP, or Either. An incoming call must then match a Connection profile in order to be accepted. If the parameter is set to Either, any authentication scheme supported by both hosts can be used, including MS-CHAP. If Recv Auth is set to None, incoming MP+ or PPP calls are not required to provide a password. (For further discussion of PAP, CHAP, and MS-CHAP, see the Chapter 7, “Setting Up Pipeline Security.”)
- 4 Set the bandwidth parameters as appropriate.
The bandwidth settings in the Answer profile apply to incoming calls for which no Connection profile exists. If a Connection profile exists, its settings take precedence. (For a discussion of bandwidth settings, see “Example MP+ configuration” on page 1-20.)
- 5 Close and save the Answer profile.

Connection profiles

Connection profiles contain parameters that define individual connections. To set up a basic Connection profile, do the following:

- 1 Open Ethernet > Connection > *any profile*.
For example:

```
20-101 Corporate-gateway
>Station=Corporate-gateway
Active=Yes
Encaps=MPP
Dial #=nnnnnnnnnnnn
Alt Dial#1=5551112
Alt Dial#2=5551113
Alt Dial#3=5551114
Calling #=
Called #=
```

Configuring WAN Connections

Connection profiles

```
Route IP=Yes
Route IPX=N/A
Bridge=N/A
Dial Brdcast=N/A
Encaps options...
Ip options...
Ipx options...
Session options...
Telco options...
```

- 2 Enter the Station name.

For example:

```
Station=Corporate-gateway
```

This is the name of the Remote end of the connection, and can be up to 31 characters.

- 3 Specify if the connection is allowed to be used or is disabled.

For example:

```
Active=Yes
```

Yes indicates the profile can be used. No deactivates the use of the connection.

- 4 Set the type of encapsulation.

For example:

```
Encaps=MPP
```

- 5 Enter the Dial #.

For example:

```
Dial # = 218005551111
```

This is the number the Pipeline dials to reach the remote network.

- 6 Enter the alternate numbers to dial. Alternate numbers are used if the dial number fails to connect. Each alternate dial number is tried before the secondary profile is used. If you do not use all three alternate number fields, leave the last field(s) blank. A blank field designates the end of the alternate numbers.

- 7 Enter the Calling #.

For example:

```
Calling # =
```

It is the number that an incoming caller's phone number is compared to, in order to authenticate the call using Id Auth.

8 Enter the Called #.

For example:

```
Called #=8005551111
```

It is usually the same as the Dial #, but any prefixes or trunk numbers are removed. The number can then be used by the far end to authenticate the call.

9 Enter whether or not the connection will Route IP, Route IPX, or Bridge unrouted protocols.

10 Enter a value for Dial brdcast.

For example:

```
Dial Brdcast=No
```

The value determines if broadcast packets initiate a connection.

11 To set the Encaps options, see “Encapsulation options” on page 1-14, “MP, MPP, and MP+ connections” on page 1-16, or “Configuring Frame Relay connections” on page 1-25, depending on the value of the Encaps parameter.

12 To set the IP options, see Chapter 2, “Configuring IP Routing.”

13 To set the IPX options, see Chapter 4, “Configuring IPX Routing.”

14 Session and Telco options are described in the following sections.

Session options

Each Connection profile contains a group of session parameters for managing WAN sessions. To set the Session options, do the following:

1 Open Ethernet > Connection > *any profile* > Session Options.

For example:

```
Session options...
>Data Filter=0
  Call Filter=0
  Filter Persistence=No
  Idle=60
  Preempt=60
  IPX SAP Filter=0
  BackUp=
```

```
Secondary=  
Block calls after=0  
Blocked duration=0
```

- 2 Set the Data and/or Call Filter parameters to prevent routine network “chatter” from keeping a connection active. (For a discussion of how to create filters, see Chapter 6, “Defining Filters and Firewalls.”)
- 3 If a filter is applied, and you want the filter to persist even if the connection is timed out or disconnected, set Filter Persistence to Yes. (For more information, refer to Chapter 6, “Defining Filters and Firewalls.”)
- 4 Set the Idle (timer) parameter to a value in seconds.
For example:

```
Idle=120
```

This specifies the Pipeline will wait 120 seconds before clearing a call when a session is inactive. If the timer expires, the Pipeline clears the call. If the parameter is set to zero, the Pipeline does not enforce a time limit.

The most common value is 120 seconds. For ISDN lines, there is often a premium charge for the first minute of any connection, so you don’t want to keep clearing and reconnecting, but you don’t want to be charged for time you don’t need. To manually clear a call, use the DO hang up command, discussed in the “DO Command Reference” chapter of the *Reference Guide*.

- 5 The Preempt parameter specifies the number of seconds the Pipeline waits before using one channel of an idle link for a new call. You can specify a number between 0 and 65535. The Pipeline sets no time limit if you enter 0 (zero). The default setting is 60.
- 6 The IPX SAP Filter is similar to the data or call filter, but prevents Netware SAP packets from unnecessarily initiating or keeping alive a connection.
- 7 The Backup and Secondary parameters are used to name other profiles that can be used if the current connection cannot be reached. The Secondary profile is used if the Backup profile is unavailable.
- 8 Block Calls After and Blocked Duration are used to control the number of times the unit redials the remote end if the line is busy, and how long to wait before retrying the connection.

For more information on any parameter, see the *Reference Guide*.

Telco options

To set the Telco options, do the following:

- 1 Open the Ethernet > Connections > *any profile* > Telco Options.

For example:

```
AnsOrig=Ans Only
Callback=No
Call Type=Switched
Group=N/A
FT1 Caller=N/A
Data Svc=56KR
Force 56=N/A
Bill #=[ ]
```

- 2 AnsOrig specifies whether the Pipeline initiates the connection, answers an incoming call, or both. Both is the default.
- 3 When Callback is Yes, the Pipeline hangs up the incoming call and calls back the remote end, using the Dial # specified in the Connection profile.
- 4 The Call Type parameter describes a type of link to a telecommunications service. Switched is the most common, since it refers to almost all connections that don't use leased or dedicated lines. (See Networking Basics on the Pipeline Companion CD for more information.)

If Nailed, Nailed/MPP, or Perm/Switched are used, you need to fill in the Group number.

A nailed connection is a permanent link that is up as long as the physical connection persists. For a nailed connection, you must specify the group number of the nailed channels. (For a discussion about group number, see “Nailed groups” on page 1-3.) For example:

```
Call Type=Nailed
Group=3
```

A nailed/MPP connection combines nailed and switched channels. If you choose this Call Type, you have to use the FT1 (fractional T1) Caller parameter to specify which side of the link can add switched channels. (See “Nailed MPP connections” on page 1-23 for additional information.)

A permanent switched connection is an outbound switched call that attempts to remain up at all times. If the unit or central switch resets or if the link is terminated, the permanent switched connection attempts to restore the link at

10-second intervals, which is similar to the way a nailed connection is maintained. A permanent switched connection causes a long connection time, but conserves connection attempts, which is cost effective for some customers. See the Call Type parameter in the *Reference Guide* for additional information.

- 5 The FT1 parameter is used only for Nailed/MPP connections to indicate which side of the connection adds bandwidth (dials up more channels).
- 6 Data Svc refers to the type of service the connection is using, such as 64 kbps, 56 Kbps, or voice. Force 56 is used to ensure that 56 Kbps is used end-to-end, even if 64 Kbps is available, since some lines in Europe and the Pacific Rim cannot use 64Kbps.
- 7 Bill # is the number (of all the numbers allocated to the service connected to the Pipeline) to which the phone company bills for this connection.

Note: For ISDN lines, this parameter is only functional in Australia.

Encapsulation options

PPP connections

A PPP connection uses PPP encapsulation on a single-channel call. To configure a PPP connection, you must perform the following tasks:

- Determine the appropriate routing, authentication, and compression settings.
- Make sure that the PPP options in the Answer profile are configured.
- Configure the PPP connection in a Connection profile.
- Configure the routing or bridging setup of the Pipeline and for the WAN connection.

Note: This section assumes that the Answer profile has been set up to enable PPP connections. (For a discussion of enabling this connection, see “The Answer profile” on page 1-7.) PPP connections are usually bridged or routed network connections initiated in PPP dialup software. (Bridging and routing configurations are discussed in their own chapters.)

Unless the Send Auth parameter is set to None, the Pipeline must be assigned a name in the Sys Config profile. To specify a name in the System menu, do the following:

- 1 Open the System > Sys Config menu.
- 2 Specify a name for the Pipeline unit in the Name parameter.
For example:
Name=MYPIPE1
- 3 Close the Sys Config menu.

To configure a PPP connection:

- 1 Open the Ethernet > Connection > *any profile* > Encaps.
For example:
Encaps=PPP
- 2 Open the Encaps Options submenu of the same profile.
Send Auth=CHAP
Send PW=*SECURE*
Recv PW=*SECURE*
MRU=1524
LQM=No
LQM Min=600
LQM Max=600
Link Comp=Stac
VJ Comp=Yes
- 3 Set the Send Auth parameter to PAP, CHAP, or MS-CHAP.
For example:
Send Auth=CHAP
Both sides of the connection must support the selected protocol. Note that MS-CHAP is only supported when both ends are using Windows NT 4.0.
- 4 Enter the password sent from the Pipeline to the remote device in the Send PW parameter's edit field. For example:
Send PW=*SECURE*
- 5 Enter the password the remote device sends to the Pipeline in the Recv PW parameter's edit field. For example:
Recv PW=*SECURE*
- 6 The values for Maximum Receive Unit (MRU) should remain at the default unless the remote cannot support it. The value defines the maximum number of bytes that can be received in a packet over PPP.

For example:.

MRU=1524

- 7 Specify if Link Quality Management (LQM) is to be used on the link, and if so, set the minimum and maximum reporting periods. Both sides of the connection must agree to use the utility.

For example:

LQM=No

LQM Min=600

LQM Max=600

- 8 If appropriate, turn on data compression.

For example:

Link Comp=Stac

VJ Comp=Yes

Press Esc to close and save the profile.

MP, MPP, and MP+ connections

MP supports multi-channel links, but not Dynamic Bandwidth Allocation (DBA). The base-channel count specifies the number of channels used for a connection. In addition, MP requires that all channels in the connection share the same phone number (that is, the channels on the answering side of the connection must be in a hunt group). MP is an extension of PPP that supports the ordering of data packets across multiple channels.

MP+ (Multilink Protocol Plus), extends the capabilities of MP (Multilink PPP) to support inverse multiplexing, session management, and bandwidth management. MP+ consists of two components: a low-level channel identification, error monitoring, and error recovery mechanism, and a session management level for supporting bandwidth modifications and diagnostics. MP+ enables the Pipeline to perform Dynamic Bandwidth Allocation (DBA)—that is, MP+ enables the Pipeline to add or remove channels without disconnecting a link as the need for bandwidth increases or decreases.

Both the dialing side and the answering side of the link must support MP+. If only one side supports MP+, the connection tries to use MP. If that fails, the connection uses standard single-channel PPP. Note that neither MP nor PPP support DBA. MPP drops the most recently connected channel first.

To configure an MP+ connection, you must perform the following tasks:

- Work with the caller to find out what networking software and Ascend Communications, Inc. configuration they have.
- Find out the required routing/bridging and authentication information for the caller.
- Configure the MP+ connection in a Connection profile.
- Configure the routing or bridging setup in the Pipeline and for the WAN connection.

Note: This assumes that the Answer profile has been set up to enable MP+ connections. (For a discussion about enabling connections, see “The Answer profile” on page 1-7.) Routing and bridging configurations are discussed in their respective chapters.

Dynamic bandwidth allocation (DBA)

DBA is part of how MP+ works, and is a way to automatically add or subtract channels on demand. When traffic levels expand, the Pipeline adds switched channels to the call. When traffic levels subside, it removes channels and frees up the bandwidth for re-allocation.

DBA uses percentage calculations to obtain average line utilization (ALU). The calculations are time sensitive. You specify a time period with the Sec History parameter and a weighting algorithm with the Dyn Alg parameter. (These parameters are set in Ethernet > Answer > PPP Options, and Ethernet > Connections > *any profile* > Encaps Options.)

When the level of activity on the line is sampled, the weight assigned to the currently required bandwidth depends on how much of the specified time period has elapsed and which weighting algorithm was selected. As shown in Figure 5-1, the weight can grow at a linear or quadratic rate or remain constant. The three Dyn Alg settings—Linear, Quadratic, and Constant—affect the ALU calculations as follows:

- Linear gives more weight to recent samples of bandwidth usage than to older samples taken during the period specified by the Sec History parameter. The weight grows at a linear rate.

- Quadratic (the default for MP+ calls) gives more weight to recent samples of bandwidth usage than to older samples taken during the period specified by the Sec History parameter. The weight grows at a quadratic rate.
- Constant gives equal weight to all samples taken during the time period.

Figure 1-1 illustrates the differences between the algorithms you can choose.




Figure 1-1. Bandwidth algorithms for MP+ calls

ALU is compared to a target percentage threshold, specified with the Target Util parameter. When ALU exceeds the threshold for a specified length of time, the Pipeline attempts to add channels. When ALU falls below the threshold for a specified length of time, the Pipeline attempts to remove channels. You specify the time periods for adding and removing channels with the Add Pers and Sub Pers parameters, respectively. (These parameters are located in Ethernet > Answer > PPP Options, and Ethernet > Connections > *any profile* > Encaps Options.)

For a discussion about removing the base channel in response to reduced bandwidth requirements, see “Clearing a call on the basis of idle bandwidth” on page 1-19. For the recommended method of bringing down inactive connections on the basis of idle time, see “Session options” on page 1-11.

Guidelines for configuring DBA

For optimum MP+ performance, both sides of a connection must have the following parameters set to the same values:

- Base Ch Count (in the Connection profile)
- Min Ch Count (in the Answer profile and the Connection profile)

- Max Ch Count (in the Answer profile and the Connection profile)

Other considerations for configuration of dynamic bandwidth allocation include:

- The values for the Sec History, Add Pers, and Sub Pers parameters should be set to ignore spikes in bandwidth utilization that last for a shorter time than it takes to add capacity.

Over T1 lines, the Pipeline can add bandwidth in less than ten seconds; over ISDN lines, the Pipeline can add bandwidth in less than five seconds.

- Once the Pipeline adds bandwidth, a minimum usage charge typically applies. Thereafter, billing is time sensitive.

The Sub Pers value should be at least equal to the time billed for the minimum duration charge plus one or two billing time increments. Typically, billing is done to the next multiple of six seconds, with a minimum charge for the first thirty seconds. Your carrier representative can help you understand the billing structure of their switched tariffs.

- Channels should not be added or removed too quickly (less than 10-20 seconds apart).

Adding or removing channels very quickly leads to many short-duration calls, each of which incurs a charge from the carrier. In addition, adding or removing channels too quickly can affect link efficiency, because the devices on either end have to retransmit data when the link speed changes.

Monitoring DBA

The DBA Monitor parameter enables you to specify which side of the link monitors traffic. Only the initiating side, however, can add or subtract bandwidth.

By default, the calling unit adds or subtracts bandwidth on the basis of how much data it transmits. To alter the default behavior, you can set the DBA Monitor parameter to Transmit-Recv, which tells the calling unit to add or subtract bandwidth on the basis of how much data it transmits *and receives*. Or, you can set it to None to tell the Pipeline not to monitor traffic over the link. If both sides of the link have DBA Monitor set to None, DBA is disabled.

Clearing a call on the basis of idle bandwidth

The Idle Pct parameter specifies a percentage of bandwidth utilization below which an MP+ call is cleared. Bandwidth utilization must fall below this

percentage on *both sides* of the connection before the Pipeline clears the call. If the device at the remote end of the link has an Idle Pct setting lower than the value you specify, the Pipeline does not clear the call until bandwidth utilization falls below the lower percentage.

The default value for Idle Pct is 0, causing the Pipeline to ignore bandwidth utilization when determining whether to clear a call. It uses the Idle timer instead.

Example MP+ configuration

For MP+ calls, you can use one authentication method for the base channel of the call, and require another password for authenticating subsequent channels as they are added. For details, see “Requesting PAP-TOKEN mode” on page 7-19.

The Pipeline must have a system name if PAP, CHAP, or MS-CHAP is to be used for outgoing calls. To assign a name to the Pipeline:

- 1 Open the System > Sys Config menu.
- 2 Specify a name for the Pipeline unit in the Name parameter.
For example:
Name=Pipe1
- 3 Close and save the Sys Config menu.

Next, configure the profile for MP+ connection:

- 1 Open Ethernet > Connection > *any profile* > Encaps.
- 2 Select MP+ encapsulation.
Encaps=MPP
- 3 Open the Encaps Options submenu of the same profile.
Send Auth=CHAP
Send PW=*SECURE*
Aux Send PW=N/A
Recv PW=*SECURE*
DBA Monitor=Transmit
Base Ch Count=1
Min Ch Count=1
Max Ch Count=2
MRU=1524

```
LQM=No
LQM Min=600
LQM Max=600
Link Comp=Stac
VJ Comp=Yes
Dyn Alg=Quadratic
Sec History=15
Add Pers=5
Sub Pers=10
Target Util=70
Idle Pct=0
Split Code.User=No
```

- 4** Specify the authentication protocol to be used.

For example:

```
Send Auth=CHAP
```

- 5** Enter the Send and Receive passwords.

For example:

```
Send PW=*SECURE*
Recv PW=*SECURE*
```

- 6** Set the number of channels the Pipeline can use for this connection.

For example:

```
Base Ch Count=1
Min Ch Count=1
Max Ch Count=2
```

- 7** If appropriate, turn on data compression.

For example:

```
Link Comp=Stac
VJ Comp=Yes
```

- 8** Configure the bandwidth options.

For example:

```
Dyn Alg=Quadratic
Sec History=15
Add Pers=5
Sub Pers=10
Target Util=70
```


Configuring WAN Connections

Connection profiles

- 9 Set the Idle Pct parameter.
For example:
`Idle Pct=0`
When this parameter is set to 0, the Idle parameter is used instead.
- 10 You can set Split Code.User to Yes so that multiple users on your LAN can use a token card to authenticate with a central server. The server must be using an Ascend RADIUS authentication server using CACHE-TOKEN-CHAP. (See more about this parameter in the *Reference Guide*.)
- 11 Close and save the Connection profile.

BACP connections

Bandwidth Allocation Control Protocol (BACP) is an industry standard that adds or removes bandwidth as needed. To use BACP, do the following:

- 1 Open Ethernet > Answer > PPP Options and set BACP to Yes.

For example:

```
PPP options...
Route IP=Yes
Route IPX=N/A
Bridge=N/A
Recv Auth=None
MRU=1524
LQM=No
LQM Min=600
LQM Max=600
Link Comp=Stac
VJ Comp=Yes
>BACP=Yes
Dyn Alg=Quadratic
Sec History=15
Add Pers=5
Sub Pers=10
```

- 2 Open Ethernet > Connection > *any profile*.
- 3 Set the Encaps Option to MP.

For example:

```
Encaps=MP
```

- 4 Open the Encaps Options submenu of the same profile and set BACP to Yes.
For example:

```
Encaps options...
  Send Auth=PAP
  Send PW=*****
  Aux Send PW=N/A
  Recv PW=N/A
  Base Ch Count=1
  Min Ch Count=1
  Max Ch Count=2
  MRU=1524
  LQM=No
  LQM Min=600
  LQM Max=600
  Link Comp=Stac
  VJ Comp=Yes
>BACP=Yes
  Dyn Alg=Quadratic
  Sec History=15
  Add Pers=5
  Sub Pers=10
  Target Util=70
```

Note: The Idle Percent parameter does not appear in the Encaps Options menu when Encaps is set to MP, as it does not apply to MP or BACP.

Nailed MPP connections

A Nailed/MPP connection is a permanent connection that can add switched channels for increased bandwidth. A Nailed/MPP connection is established when its nailed or switched channels are connected end-to-end.

Switched channels are added to or subtracted from the Nailed/MPP connection as required by the DBA parameters of either the far-end or near-end Connection profile. If the two sides of a connection disagree on the number of channels needed for a connection, the side requesting the greater number prevails. Calculations on the required number of channels are made by each side based on the traffic received at that side.

The maximum number of channels for the Nailed/MPP connection is either the Max Ch Count or the number of nailed channels in the specified group, whichever is greater. If a nailed channel fails, the Pipeline replaces that channel with a switched channel, even if the call is online with more than the minimum number of channels.

To configure a nailed MPP connection, first configure a regular MP+ connection (see “Example MP+ configuration” on page 1-20). Then follow these steps:

- 1 Open the Telco Options submenu of the Connection profile.
- 2 Specify the Nailed/Mpp call type.
For example:
`Call Type=Nailed/Mpp`
- 3 Specify the group number of the nailed channels. (For a discussion of group numbers, see “Nailed groups” on page 1-3.)
For example:
`Group=1, 2`
- 4 Specify that the Pipeline is the designated caller for the switched part of the connection.
For example:
`AnsOrig=Call Only`
`FT1 Caller=Yes`
- 5 Close and save the profile.

On the far end of the connection, set the AnsOrig and FT1 Caller parameters for answering only. Note that the DO Hangup command only works from the caller end of the connection.

You can reconfigure the parameters of a Nailed/MPP Connection profile at any time, but the changes become active only after the call is brought down and then back up. However, if you add a value to the Group parameter and save the change, the additional channels are added to the connection without having to bring it down and back up. For example, changing from Group=1 to Group= 2 as described in “Nailed groups” on page 1-3.

Note: If a Nailed/MPP connection is down and the nailed channels are also down, the connection does not reestablish itself until the nailed channels are brought back up or the switched channels are dialed. (The switched channels are

dialed when the calling unit receives a packet whose destination is the unit at the far-end of the Nailed/MPP connection.)

Configuring Frame Relay connections

Some Ascend units provide Frame Relay as an option. If you are not sure if your unit supports Frame Relay, press the Tab key to highlight the Sys Option status window and then use the arrow key (or press Ctrl-N) to scroll down in the window. If Frame Relay is installed, this text appears in the status window:

Frm Rel Installed

Frame Relay profiles define connections between the Pipeline and Frame Relay switches. The connections are almost always nailed. Switched connections can be used only in the rare situation in which the Frame Relay network allows dial-in connections, and connections to the network are always initiated by the Pipeline. (Frame relay switches currently have no dial-out connection capability.)

Connection profiles define logical links to an end-point on the Frame Relay network. Each Connection profile must specify a Data Link Connection Identifier (DLCI) for that link. A DLCI is a number between 16 and 991, which is assigned by the Frame Relay administrator. A DLCI is not an address, but a local label that identifies a logical link between a device and a Frame Relay switch. (That is, the DLCIs enable the Frame Relay switch to identify the logical link associated with each Connection profile.) The switch uses the DLCI to route frames through the network, and the DLCI may change as frames are passed through multiple switches.

Note: You need at least one Frame Relay profile and Connection profile to define a logical link to the Frame Relay network.

To configure a Frame Relay connection, you must perform the following tasks:

- Make sure that nailed channels are available for the link to the Frame Relay switch.
- Configure a Frame Relay profile that uses those channels to connect to the Frame Relay (FR) switch.
- Obtain the DLCIs you need from the Frame Relay administrator (at the telephone company, or your network administrator).

Configuring WAN Connections

Configuring Frame Relay connections

Each connection requires its own DLCI.

- Obtain the routing/bridging information for the remote network.
- Make sure that the Answer profile enables FR encapsulation.
- Configure the Frame Relay connection in a Connection profile.
- Configure the routing or bridging setup in the Pipeline and across the WAN connection.

Note: This section focuses on configuring Frame Relay connections. It assumes that the Answer profile has been set up to enable such connections. (For a discussion about enabling connections, see “The Answer profile” on page 1-7.) Configuring bridging and routing is discussed in subsequent chapters.

Example of options used to configure logical links

A Connection profile defines a logical link to an end-point reached through a Frame Relay switch. The Pipeline supports Frame Relay “Gateway” mode. A Frame Relay gateway connection is a bridging or routing link between the Pipeline and a remote network via a Frame Relay switch. When the Pipeline receives IP packets destined for that network, it encapsulates the packets in Frame Relay (as specified in RFC 1490) and forwards the data stream with the specified DLCI to the Frame Relay switch. The Frame Relay switch uses the DLCI to route the frames to the right destination.

Figure 1-2 shows a Pipeline with three gateway connections to customer premise equipment (CPE) at remote sites across the Frame Relay network. Gateway

connections can support bridging and routing, so the Pipeline can forward any type of protocol traffic from the local network onto the Frame Relay network.

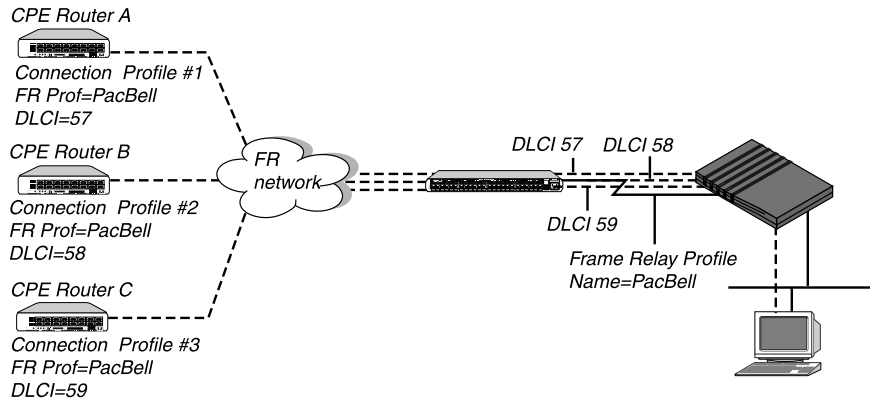


Figure 1-2. Gateway connections to the Frame Relay network

Connection profiles #1, #2, and #3 use Frame Relay encapsulation (RFC 1490) and include both a DLCI number for the logical link and the name of the Frame Relay profile for the nailed connection. The Frame Relay profile defines a nailed connection between the Pipeline and a Frame Relay switch. The Connection profiles and the Frame Relay profile in this example are defined below:

Table 1-1. Frame Relay and gateway profiles

Connection profiles (gateway)	Frame Relay profile
Station=CPEA Active=Yes Encaps=FR Encaps options... FR Prof=PacBell DLCI=57	Name=PacBell Active=Yes Call Type=Nailed Nailed Grp=1 Data Svc=64K Link Mgmt=T1.617D...

Table 1-1. Frame Relay and gateway profiles (continued)

Connection profiles (gateway)	Frame Relay profile
Station=CPEB Active=Yes Encaps=FR Encaps options... FR Prof=PacBell DLCI=58	See profile above.
Station=CPEC Active=Yes Encaps=FR Encaps options... FR Prof=PacBell DLCI=59	See profile above.

Configuring a Frame Relay profile

To define the Frame Relay profile:

- 1 Open Ethernet > Frame Relay > *profile*.

For example:

```
Name=PacBell17
Active=Yes
Call Type=Nailed
FR Type=DTE
LinkUp=No
Nailed Grp=1
Data Svc=64k
Dial #=N/A
Link Mgmt=T1.617D
N391=6
DTE N392=3
DTE N393=4
DCE N392=
DCE N393=
T391=10
```

T392=15
MRU=1532

- 2 Assign the profile a name.

For example:

Name=PacBell

The name can contain up to 15 alphanumeric characters. You have to use this name in Connection profiles that use this connection to the switch.

- 3 Activate the profile.

For example:

Active=Yes

- 4 Specify that this is a nailed connection.

For example:

Call Type=Nailed

- 5 Specify the Frame type of service.

For example:

FR Type=DTE

Your service provider will give you this information.

- 6 Specify whether the Frame Relay link comes up automatically and stays up, even when the last DLCI has been removed.

For example:

LinkUp=No

- 7 Enter the group number of the nailed channels to be used.

For example:

Nailed Grp=1

Nailed is the default for Frame Relay connections. When the call type is nailed, dial numbers and other telephone company parameters are N/A. You can specify switched if the Frame Relay switch allows dialing in. However, Frame Relay networks currently have no dial-out connection capability.

- 8 Set the data service.

For example:

Data Svc=64k

- 9 Specify the link management protocol used between the Pipeline and the Frame Relay switch.

For example:

Link Mgmt=T1.617D

If you specify Link Mgmt=T1.617D, set the following additional parameters:

N391

DTE N392

DTE N393

T391

T392

N391 specifies how many polling cycles the Pipeline waits before requesting a full status report. DTE N392 is the maximum number of error events that can occur in the sliding window defined by DTE N393. DTE N393 specifies the width of the sliding window used by the DTE N392 parameter.

T391 specifies the number of seconds between Status Enquiry messages.

T392 specifies the number of seconds that the Pipeline waits for a Status Enquiry message before recording an error.

See the *Reference Guide* for more details.

- 10 Close the Frame Relay profile.

Configuring a gateway connection

This section shows how to configure a Frame Relay gateway connection. Routing and bridging parameters must also be configured to have a working connection.

To configure a Frame Relay gateway connection to Customer Premises Equipment (CPE) on the Frame Relay network:

- 1 Open Ethernet > Connection > *any profile*.

- 2 Specify the name of the CPE.

For example:

Station=CPEA

- 3 Activate the profile.

For example:

Active=Yes

- 4 Select Frame Relay encapsulation.

For example:

```
Encaps=FR
```

The Pipeline uses this encapsulation method to encapsulate packets before routing them out to the CPE, and removes the Frame Relay encapsulation from packets coming in from the CPE.

- 5 Open the Encaps Options submenu of the same profile.

```
FR Prof=Pac Bell
```

```
DLCI=17
```

- 6 Set the DLCI parameter to the number assigned by the Frame Relay administrator.

For example:

```
DLCI=500
```

The Frame Relay administrator must assign the DLCI number. It determines how packets will be routed at the Frame Relay switch.

- 7 Specify the name of the Frame Relay profile that defines the nailed connection to the Frame Relay switch.

For example:

```
FR Prof=PacBell
```

The name must match the Name parameter in the Frame Relay profile exactly, including case changes.

- 8 Close and save the profile.

Inverse ARP for Frame Relay

Inverse Address Resolution Protocol (InARP) allows a device to resolve the protocol address of another device when the hardware address is known. In the case of Frame Relay the hardware address is the DLCI. The Ascend implementation of Inverse ARP responds only to Frame Relay and IP Inverse ARP requests.

The ARP protocol type for Inverse ARP requests must be IP(0x8000). ARP hardware address type must be the 2-byte Q.922 address. All other types are discarded.

The Inverse ARP response supplies the following data:

- ARP source protocol address is the IP address of the Pipeline, found in the Ethernet > Mod Config > Ether Options > IP Adrs parameter.
- ARP source hardware address is the Q.922 address of the local DLCI.

Note: The Pipeline does not issue any Inverse ARP requests.

Refer to RFCs 1293 and 1490 for details on Inverse ARP.

Configuring a T1 line

The T1 line is the RJ-45 port on the back panel of the Pipeline unit that has the tab hole facing up. This line is not channelized, but you can configure it to act like a channelized T1 with a number of DS0 channels as specified by your carrier.

The nailed T1 line also requires some port configuration; for example, you must specify the signals that indicate that the DCE (Data Communication Equipment) is ready to connect. In addition, you may need to set the amount of attenuation that the Pipeline should apply to the line's network interface in order to match the cable length from the Pipeline to the next repeater.

Configuration steps

To configure the nailed T1 line, you will perform the following tasks:

- Specify a group number associated with the nailed T1 line.
You assign a group number to the line and then specify that group number in Connection profiles that will access the WAN across this interface.
- Activate the port.
- Supply carrier information, such as encoding, framing, and buildout (attenuation).

To configure the nailed T1 line:

- 1 Open the Nailed T1 > Mod Config menu.

```
Nailed T1 Group=3
Activation=Enabled
Framing Mode=D4
```

```
Encoding=B8ZS
FDL=None
First DSO=1
Last DSO=24
Clock Source=No
Loopback=Normal
```

- 2** Set Nailed T1 Group to a unique value.

For example:

```
Nailed T1 Group=3
```

This parameter must specify a group number that has *not* already been assigned to channels on another line.

- 3** Activate the nailed T1 line.

```
Activation=Enabled
```

- 4** Set the T1 framing mode.

For example:

```
Framing Mode=D4
```

- 5** Set the Encoding parameter as specified by your carrier.

For example:

```
Encoding=B8ZS
```

Encoding refers to the way in which data is represented by the digital signals on the line. Both sender and receiver must agree on the type of encoding in use in order to accurately interpret the value of a signal. B8ZS is often required for ISDN.

- 6** Set the Facilities Data Link (FDL) protocol to enable Channel Service Unit (CSU) capability on the line.

For example:

```
FDL=ANSI
```

FDL is a 4 Kbps system-data channel only available when using Extended Super Frame (ESF) format. FDL provides information at regular intervals to the carrier's maintenance devices, enabling the telco to use the FDL protocol to monitor the quality and performance of T1 lines. Your carrier can tell you which FDL protocol to specify.

- 7** Enter the first DS0 channel assigned to this line by your carrier.

For example:

Configuring WAN Connections

Configuring a T1 line

First DS0=1

- 8 Enter the last DS0 channel assigned to this line by your carrier.

For example:

Last DS0=24

You must enter a number equal to or greater than the first DS0 before you can save the profile.

Note: If you selected Alternate Mark Inversion (AMI) for encoding, you cannot have more than six DS0 channels. B8ZS allows a maximum of 24 channels.

- 9 Set the clock source. For example:

Clock Source=No

No assumes that the telco provides the clock source. If you are connecting two Pipeline 130s back-to-back, set Clock Source on one Pipeline to Yes to invoke the unit's internal oscillator as the clocking source, and set Clock Source on the other Pipeline to No.

- 10 Set Loop Back mode.

For example:

Loop Back=Normal

When set to Normal, communication can occur. To put the unit in loopback mode, set this value to Relay Loopback, Line Loopback, or Data Loopback, depending on the type of test you need.

Note: Once you enable a loopback mode and save the profile, no communication is possible over the WAN.

- 11 Set the attenuation (Buildout) if appropriate.

For example:

Buildout=0db

If you specify a value other than 0 decibels (the default) for Buildout, the Pipeline applies an attenuator to the T1 line, causing the line to lose power. (Repeaters boost the signal on a T1 line—if the Pipeline is too close to a repeater, you need to add some attenuation.)

- 12 Close and save the profile.

Backup Connection disconnect timer

When the nailed T1 connection on the Pipeline fails, the ISDN backup connection is established, which allows the ISDN backup connection to be timed out when the nailed T1 connection is reestablished.

The Pipeline provides a backup ISDN connection for its primary nailed T1 line. The Pipeline senses when the primary connection has been reestablished and routes all traffic through this primary connection. At this point, the ISDN line's idle timer is activated, and the ISDN link is automatically timed out.

Configuring a V.35 serial WAN port

The serial WAN data rate is determined by the clock speed received from the link. The maximum acceptable clock is 1.56Mbps. The clock speed at the serial WAN port has no effect on the bandwidth of other WAN interfaces in the Pipeline.

When the V.35 serial WAN port is enabled, all Connection profiles are sampled once every 10 seconds. If a Connection profile or Frame Relay profile is configured for leased-line operation and the group parameter in the Connection profile or a Frame Relay profile is set to the same value as the Group parameter in the V.35 Mod Config profile, the V.35 port is set for synchronous HDLC mode and an attempt is made to bring up the connection on that port.

Before configuring the Serial WAN profile, decide which Frame Relay or other connection will use the V.35 Serial WAN port. Then set the group parameter for that profile to the same value as the Group parameter in the V.35 Mod Config profile:

- For Connection profiles, set the Group parameter in the Ethernet > Connections > *profile* > Telco Options submenu
- For a Frame Relay profile set the Nailed Grp parameter.

To configure the V.35 Serial WAN port:

- 1** Open the Serial WAN profile.

```
30-000 Serial WAN
Mod Config...
```

Configuring WAN Connections

Configuring a V.35 serial WAN port

- 2 Open the Mod Config profile.
20-B00 Mod Config...
Module Enabled=Yes
Group=3
Activation=Static
- 3 Set Module Enabled to Yes.
- 4 Select a Group number that matches the Group parameter in a Connection profile or the Nailed Grp parameter in a Frame Relay profile.
- 5 Select the appropriate Activation.
This selects the signals at the serial WAN port that indicate that the Data Circuit-Terminating Equipment (DCE) is ready to connect.
- 6 Close the Serial WAN profile and save the changes.

Configuring IP Routing

This chapter contains the following topics:

Introduction to IP routing on the Pipeline	2-1
Managing the routing table.	2-14
Configuring IP routing connections	2-30
Ascend Tunnel Management Protocol (ATMP)	2-37

Introduction to IP routing on the Pipeline

An IP router moves data towards its destination using the most efficient path it knows. IP routers keep track of the source and destination addresses of packets it handles, builds tables with this information, collects information in routing tables of other routers, and can advertise its own routes. (For information about routing packets using the Internet Packet eXchange protocol used in NetWare LANs, see Chapter 4, “Configuring IPX Routing.”)

The most common uses for IP routing connections in the Pipeline are to:

- Enable IP connections to the Internet (through Internet Service Providers).
- Connect distributed IP subnets to a corporate backbone (telecommuting and remote office hubs).

The Pipeline supports IP routing over PPP, MP, MP+, and Frame Relay connections. The Pipeline is fully interoperable with non-Ascend products that conform to the TCP/IP protocol suite and associated RFCs.

IP routing connections have a level of built-in authentication, because the Pipeline matches the IP address of a Connection profile to the source IP address of a caller. For most sites, however, this level of security is not enough and a form of password authentication is used as well. (For more information, see Chapter 7, “Setting Up Pipeline Security.”)

Note: IP routing can be configured along with protocol-independent bridging and IPX routing in any combination. However, you cannot bridge and route IP packets across the same connection. When you configure the Pipeline as an IP router, IP packets are no longer bridged at the link layer. They are *always* routed at the network layer. All other protocols continue to be bridged unless you turn off bridging. (For more information about bridging, see Chapter 5, “Configuring the Pipeline as a Bridge.”)

Host-to-router connections

When the device connecting to the Pipeline is a host running PPP dial-in software, the Pipeline adds a “host route” to its routing table. (For discussion of host routes, see “Subnet mask notation” on page 2-4.)

If the host belongs to its own IP network, the Pipeline must have a Connection profile stating the host’s address, using a 32-bit netmask.

When the dial-in user calls the Pipeline, the Pipeline checks its Answer profile to verify that it can accept incoming IP routing calls. If it can, it checks whether it has a Connection profile for that user. If the Answer profile is not configured to allow incoming IP routing calls, or if there is no Connection profile for the call, the call is refused.

If the Answer profile does allow the call, the Pipeline looks for a Connection profile matching the user’s name and IP address. If the Pipeline doesn’t find a matching profile, it ends the call.

If the Pipeline does find an address and the PPP software accepts it, the Pipeline authenticates the connection using password authentication, and then establishes the connection.

When the connection is established, the Pipeline adds a host route to its routing table and begins functioning as an IP router between its local and WAN

interfaces. If the Pipeline is configured for RIP, it also broadcasts its updated routing table to other hosts.

Router-to-router connections

When the device connecting to the Pipeline is an IP router that belongs to an IP network, the connection results in a route to that remote network (or subnet). For example, Figure 2-1 shows a Pipeline connected to a remote router. The two Ethernet segments are separate IP networks.




Figure 2-1. An IP routing connection between two networks

If a user (attached to the Pipeline) launches a Telnet session and enters an address at a remote site (such as one on the other side of the CPE router), the remote router receives the outbound TCP/IP packets and consults its routing table. If it does not find a route to the target site, it either forwards the packets to its default router or drops them, depending on how it is configured. If it finds a route to the target site, it opens the appropriate profile and dials out.

When receiving a call, the Pipeline checks its Answer profile to verify that it can accept incoming IP routing calls. Then it does the following:

- If the Answer profile does not have IP routing on, the Pipeline ends the call.
- If the Answer profile does have IP routing on, the Pipeline looks for a profile that matches the IP address offered during PPP negotiation.
- If the Pipeline doesn't find a matching Connection profile, it ends the call.
- If the Pipeline finds a matching profile, it authenticates the connection.

After a connection is established, the Pipeline adds a network route to its routing table and begins functioning as an IP router between its local and WAN interface. If the Pipeline is configured for RIP, it also broadcasts its updated routing table to other hosts.

Subnet mask notation

In the Pipeline, IP addresses are specified in decimal format (not hexadecimal).
For example:

198 . 5 . 248 . 40

If no netmask is specified, the Pipeline assumes a default netmask based on the “class” of the address:

Table 2-1. IP address classes and default netmasks

Class	Address range	Network bits
Class A	0.0.0.0 → 127.255.255.255	8
Class B	128.0.0.0 → 191.255.255.255	16
Class C	192.0.0.0 → 223.255.255.255	24
Class D	224.0.0.0 → 239.255.255.255	N/A
Class E (reserved)	240.0.0.0 → 247.255.255.255	N/A

For example, a class C address such as 198.5.248.40 has 24 network bits, as shown in Figure 2-2. That leaves 8 bits for the host portion of the address, so up to 255 hosts can be supported on the class C network.




Figure 2-2. A class C address

To specify a netmask, the Pipeline does not use dotted decimal format, as in:

IP Address=198 . 5 . 248 . 40
Netmask=255 . 255 . 255 . 248

Instead, it includes a netmask modifier that specifies the total number of network bits in the address. For example:

198.5.248.40/29

In the example address shown above, the /29 specification indicates that an additional 5 bits of the address will be interpreted as a subnet number.

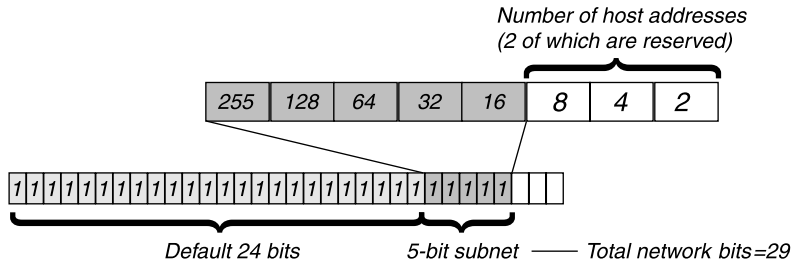


Figure 2-3. A 29-bit netmask and number of supported hosts

Eight bit-combinations are possible in 3 bits. Of those eight possible host addresses, two are reserved:

- 000 — Reserved for the network base (the cable)
- 001
- 010
- 100
- 110
- 101
- 011
- 111 — Reserved for the broadcast address of the subnet

Table 2-2 shows how standard subnet address format relates to Ascend notation for a class C network number.

Table 2-2. Standard netmasks and Ascend netmask notation

Netmask	Ascend notation	Number of host addresses
255.255.255.0	/24	254 hosts + 1 broadcast, 1 network base
255.255.255.128	/25	126 hosts + 1 broadcast, 1 network base
255.255.255.192	/26	62 hosts + 1 broadcast, 1 network base

Configuring IP Routing

Introduction to IP routing on the Pipeline

Table 2-2. Standard netmasks and Ascend netmask notation (continued)

Netmask	Ascend notation	Number of host addresses
255.255.255.224	/27	30 hosts + 1 broadcast, 1 network base
255.255.255.240	/28	14 hosts + 1 broadcast, 1 network base
255.255.255.248	/29	6 hosts + 1 broadcast, 1 network base
255.255.255.252	/30	2 hosts + 1 broadcast, 1 network base
255.255.255.254	/31	invalid netmask (no hosts)
255.255.255.255	/32	1 host — a host route

Note: A host route is a special case IP address with a subnet mask of /32; for example, 198.5.248.40/32. Host routes are required for a dial-in host.

The broadcast address of any subnet is always all ones. The network base address represents the network cable itself, which is always address 0. For example, if the Pipeline configuration assigns the following address to a remote Pipeline router:

198.5.248.120/29

the Ethernet attached to that router has the following address range:

198.5.248.120 — 198.5.248.127

The “0” address (198.5.248.120) is reserved for the cable itself. The broadcast address is 198.5.248.127, and the router itself uses one of the host addresses. That leaves five remaining host addresses on that remote subnet, which can be assigned in any order to five hosts on that subnet.

As another example, if the Pipeline configuration assigns the following address to a remote router:

192.168.8.64/26

the Ethernet attached to that router has the following address range:

192.168.8.64 — 192.168.8.127

The “0” address for this subnet is 192.168.8.64. The broadcast address must be the network base address plus six ones (six ones in base 2 equals 63 decimal, and $64+63=127$) 192.168.8.127.

Note: Early implementations of TCP/IP did not allow zero subnets. That is, subnets could have the same base address that a class A, B, or C network would have. For example, the subnet 192.168.8.0/30 was illegal because it had the same base address as the class C network 192.168.8.0/24, while 192.168.8.4/30 was legal. (192.168.8.0/30 is called a zero subnet, because like a class C base address, its last octet is zero.) Modern implementations of TCP/IP allow subnets to have base addresses that might be identical to the class A, B, or C base addresses. Ascend's implementations of RIP 2 treats these so-called zero subnetworks the same as any other network. However, it is important that you treat zero subnets consistently throughout your network. Otherwise, you will encounter routing problems.

IP routing in the Answer profile

Before the Pipeline answers an incoming call, it checks the settings in its Answer profile for information about what to do. If the call does not include the information required by the Answer profile, the Pipeline hangs up.

The parameters listed below are related to IP routing in the Answer profile. (For detailed information about each parameter, see the *Reference Guide*. You might also want to refer to other sections in this guide, including, “Configuring RIP for incoming WAN connections” on page 2-22, for setting the RIP parameter in the Answer profile, and for authentication, Chapter 7, “Setting Up Pipeline Security.”)

- Ethernet > Answer > Session Options
RIP=Off
- Ethernet > Answer > PPP Options
Route IP=Yes
Recv Auth=Either

To enable the Pipeline to answer incoming IP routing calls:

- 1 Open the Ethernet > Answer > PPP Options menu.
- 2 Turn on IP Routing.

Route IP=Yes

3 Set Recv Auth=Either.

Or set it to PAP, CHAP, or MS-CHAP. Either indicates any protocol that both sides agree upon.

Connection profiles and IP routes

The Pipeline creates a routing table when it powers up. It adds all known routes to the table, including connected routes (such as Ethernet) and routes configured in its resident Connection profiles and Static Rtes profiles. If RIP is enabled in the Ethernet, it supplies information about routes learned from other routers to the routing table. If RIP is enabled on an active connection, it supplies information about the routes received from the far-end of that connection to the routing table.

There are some static routes that the Pipeline cannot read at power-up. They do not become part of the routing table until they are up and usable. Such routes include those added via the `Iproute add` terminal server command.

How the Pipeline uses its routing table

When the Pipeline receives an IP packet whose destination address is not on the local network, it checks its routing table for the destination network and:

- If it finds a route to that network, it forwards the packet to the gateway indicated by that route. If the gateway is not local, the Pipeline opens a WAN connection to forward the packet.
- If it does not find a route to that network, it forwards the packet to the default router.
- If it does not find a route to that network and no default route has been configured, it drops the packet.

When the Pipeline receives an incoming IP routing call, it examines the source IP address and looks for a matching profile. If the source matches a resident Connection profile, the Pipeline updates its routing table, if necessary, with the route to the source network.

If the Answer profile is configured without authentication requirements (an unlikely scenario) and Profile Req'd is set to No, the Pipeline accepts any IP

routing connection that comes in. In that case, it does not have a route for the incoming source IP address, and builds a temporary route using an assumed Class A (8), B (16), or C (24) netmask for the source IP address. If this type of connection is with a router or a host that does not recognize the initial temporary route (that is, one from another manufacturer), you might have to turn on RIP or configure a static route to build a route to that network.

RIP-v2 and RIP-v1 routing

The Pipeline includes a Routing Information Protocol (RIP) version 2 implementation (RIP-v2), which includes a set of improvements to RIP-v1. You can configure the Pipeline to send, receive, or send and receive RIP-v1 or RIP-v2 on Ethernet or any WAN interface.

Note: RIP-v2 is a compatible upgrade to RIP-v1, but do not run RIP-v2 and RIP-v1 on the same network in such a way that the routers receive each other's advertisements. RIP-v1 "guesses" subnet masks, while RIP-v2 handles them explicitly. Running the two versions on the same network can result in RIP-v1 "guesses" overriding accurate subnet information obtained via RIP-v2.

RIP-v2 includes the following improvements to RIP-v1:

Subnet routing	The biggest difference between RIP-v1 and RIP-v2 is the inclusion of subnet mask information in RIP-v2 routes. RIP-v1 recognized subnet information only within the subnet and purposely did not advertise subnet masks to other routers. There was no way to distinguish between a subnet and a host entry, unless it was for a router directly connected to the subnet. When a RIP-v1 router receives an IP address, it assumes the default subnet mask. RIP-v2 passes the netmask in parallel with the address. This enables support not only of reliable subnet routing, but also of variable length masks within the same network as well as Classless Inter-domain Routing (CIDR). If a RIP-v1 router receives a RIP-v2 update that includes netmasks, it ignores the subnet information.
----------------	---

Authentication	<p>RIP-v1 provided no way of authenticating its routing advertisements. Any program that transmitted packets on UDP port 520 was considered a router with valid distance vectors.</p> <p>RIP-v2 packets include an authentication field that can contain a simple password. If a RIP-v1 router receives a RIP-v2 packet that contains a password, it ignores the field.</p>
Routing domains	<p>To enable multiple networks to share a common backbone, RIP-v2 uses a routing domain number that enables routers to recognize packets bound for a particular domain number in the router's networks.</p>
Multicasting	<p>RIP-v1 uses a broadcast address for sending updates, so its tables are received not only by routers but by all hosts on the cable as well.</p> <p>RIP-v2 uses an IP multicast address or MAC address for periodic multicasts to RIP-v2 routers.</p>

Interface-based routing

All Pipeline units implement what is referred to as system-based or box-based routing. With system-based routing, the entire box is addressed with a single IP address. For systems that have a single backbone connection, system-based routing is by far the simplest form of routing from both a configuration and trouble-shooting perspective. The alternative form of routing is referred to as interface-based routing. With interface-based routing, each physical or logical interface on the box has its own IP address.

However, there are some applications that the Pipeline is used for in which it might be useful to number some of the interfaces—in other words, to have the Pipeline operate as a partially system-based router and partially interface-based router. Reasons for using numbered interfaces include troubleshooting leased point-to-point connections and forcing routing decisions between two links going to the same final destination. More generally, interface-based routing allows the Pipeline to operate more nearly the way a multi-homed Internet host behaves, if that is needed.

Interfaced-based routing lets you configure each link as numbered (interface-based) or unnumbered (system-based). If no interfaces are specified as

numbered, then the unit operates exactly as it does when using unnumbered routing. Configure interface numbering in the Connection profile.

System behavior with a numbered interface

If a Pipeline is using a numbered interface, the following differences in operation should be noted, compared to unnumbered (system-based) routing:

- IP packets generated in the Pipeline and sent to the remote address use an IP source address corresponding to the numbered interface, not to the default (Ethernet) address of the Pipeline.
- During authentication of a call placed from a Pipeline using a numbered interface, the Pipeline reports the address of the interface as its IP address.
- The Pipeline adds, as host routes to its routing table, all numbered interfaces listed in Connection profiles.
- The Pipeline accepts IP packets whose destination are a numbered interface listed in a Connection profile, considering them to be destined for the Pipeline itself. (The packet might actually arrive over any interface, and the numbered interface corresponding to the packet's destination address need not be in the active state.)

Configuring interface-based routing

Configure interface-based routing in the IP Options submenu of the Connection profile. The IF Adrs parameter specifies the IP address of the interface. If you leave the field at its default value (0.0.0.0/0), the interface is unnumbered.

The profile below shows settings for a numbered interface. The WAN Alias parameter contains the address of the remote end, and the IF Adrs parameter contains the interface number of the near end.

```
Ip options...
  LAN Adrs=192.168.6.29/24
  WAN Alias=192.1.1.17
  IF Adrs=192.1.1.8/30
  Metric=0
  Preference=2
  Private=No
  RIP=Off
  Pool=0
```

Specifying the remote interface address

This section provides some guidelines on using interface-based routing.

If both the system and interface addresses are known

If you are adding interface-based routing to a system set up for system-based routing, enter the remote interface address in the WAN Alias parameter of the Connection profile. WAN Alias identifies the remote end of the link. If a WAN Alias is set, the following processes occur:

- Host routes are created to LAN Adrs and WAN Alias, and the WAN Alias is listed in the routing table as a gateway (next hop) to the Lan Adrs.
- A route is created to the remote system's subnet, showing the WAN Alias as the next hop.
- Incoming PPP/MPP calls must report their IP addresses as the WAN Alias (rather than the Lan Adrs). That is, the caller must be using a numbered interface, and its interface address must agree with the WAN Alias on the receiving side.

To create static routes to hosts at the remote end, use the WAN Alias address as the “next hop” (gateway) field. (The Lan Adrs address will also work, as it is for system-based routing.)

If only the interface address is known

You can omit the remote side's system address from the profile and use interface-based routing exclusively. This is an appropriate mechanism if, for example, the remote system is on a backbone net which might be periodically reconfigured by its administrators, and you want to refer to the remote system only by its mutually agreed-upon interface address.

In this case, the remote interface address is entered in the Lan Adrs parameter, and the WAN Alias is left as default (0.0.0.0). Note that Lan Adrs must always be filled in, so if the only known address is the interface address, it must be placed in the Lan Adrs parameter rather than the WAN Alias parameter.

If the remote interface address is placed in the Lan Adrs parameter, the following will take place:

- A host route is created to the Lan Adrs (interface) address.
- A net route is created to the subnet of the remote interface.
- Incoming PPP/MPP calls must report their IP addresses as the Lan Adrs (interface) address.

If the remote interface address is not specified

If interface-based routing is in use and the local interface is numbered, the remote address will usually be known (in practice, the subnet must be agreed upon by administrators of both sites). It is possible, but not recommended, to number the local interface, omitting the interface address of the remote site and using only its system or LAN address. In that case, do not use the (supposedly unknown) remote interface address in any static routes.

When a local interface is numbered but no corresponding remote interface address is set, the remote interface must have an address on the same subnet as the local, numbered interface. Incoming PPP will be rejected if the Connection Profile numbers the local interface and the (remote) caller supplies an address not on the same subnet.

Multicast forwarding and IGMP functionality

The Pipeline supports Internet Group Membership Protocol (IGMP) version 1 and version 2, that enable the Pipeline to subscribe as a multicast client. The Pipeline transparently passes any multicast traffic it receives from a multicast router to its Ethernet, making these packets available to local hosts on its Ethernet which have been set up to listen to them. The Pipeline does not operate as an IGMP router, and does not forward multicast packets to IGMP clients.

To enable multicast forwarding:

- 1 Open Ethernet > Mod Config > Multicast menu.
- 2 Set Multicast Forwarding to Yes.
This setting enables multicast forwarding in the Pipeline. The Pipeline then receives Internet Group Membership Protocol (IGMP) queries from the router and responds to them using IGMP.
- 3 Specify the name of a resident profile to be defined as the Multicast Profile.
This profile is used to connect over the WAN to the multicast router in the

IP-only version of the Pipeline. If no profile name is specified and Multicast Forwarding is set to Yes, the Pipeline assumes that its Ethernet is the Multicast interface.

- 4 Reset the Pipeline for the changes to take effect.

Managing the routing table

The Pipeline routing table is created when the Pipeline powers up. (Which routes are included and when is discussed in “Connection profiles and IP routes” on page 2-8.) To manage the routing table, you can perform one or more of the following tasks:

- Configure static routes in the IP Options of a Connection profile.
- Configure a default route for packets with an unknown destination.
- Turn off ICMP Redirects.
- Configure RIP-v1 or RIP-v2 on Ethernet.
- Turn off RIP on WAN connections.
- Assign a preference for RIP or static routes (known as route preferences).
- Display the routing table.

Parameters that affect the routing table

The list below shows parameters that affect the Pipeline IP routing table:

- Ethernet > Mod Config
 - RIP Policy=Poison Rvrs (RIP-v1 only)
 - RIP Summary=Yes (RIP-v1 only)
 - ICMP Redirects=Accept
 - Adv Dialout Routes=Trunks Up

Note: When more than one Pipeline is in use in redundant configurations on the same network, you can use the Adv Dialout Routes parameter to instruct the Pipeline to stop advertising IP routes that use dial services if its trunks are in the alarm condition. If a redundant Pipeline loses its dialout lines temporarily, and the Adv Dialout Routes parameter is set to Always, that unit continues to receive outbound packets that should be forwarded to

the redundant Pipeline. To prevent the problem, set Adv Dialout Routes to Trunks Up. For details on these parameters, see the *Reference Guide*.

- Ethernet > Mod Config > Ether Options
IP Adrs=10.2.3.2/245
2nd Adrs=0.0.0.0/0
RIP=Both-v2
RIP2 Use Multicast=Yes
Ignore Def Rt=No
- Ethernet > Connections > *any profile*
Route IP=Yes
- Ethernet > Connections > *any profile* > IP Options
LAN Adrs=10.9.8.10/22
WAN Alias=0.0.0.0
Metric=1
Preference=100
Private=No
RIP=Off
- Ethernet > Static Rtes > *any profile*
Name=SITEBGW
Active=Yes
Dest=10.2.3.0/24
Gateway=10.2.3.4
Metric=2
Preference=100
Private=No
- Ethernet > Answer > PPP Options
Route IP=Yes
- Ethernet > Answer > Session Options
RIP=Both-v2

For details about each parameter, see the *Reference Guide*.

Static and dynamic routes

A static route is a path from one network to another, which specifies the destination network and the router to use to get to that network. For routes that must be reliable, the administrator often configures more than one path (adds a secondary route), in which case the Pipeline chooses the primary route on the basis of an assigned metric.

A dynamic route is a path to another network that is “learned” dynamically rather than configured in a profile. A router that uses RIP broadcasts its entire routing table every 30 seconds, updating other routers about which routes are usable. Hosts that run ICMP can also send ICMP Redirects to offer a better path to a destination network.

Note: A dynamic route can overwrite or “hide” a static route to the same network if the dynamic route’s metric is lower than that of the static route. However, dynamic routes age and if no updates are received, they eventually expire. In that case, the “hidden” static route reasserts itself and is reinstated in the routing table.

Configuring static routes

Every Connection profile that specifies an explicit IP address is a static route. (For details on configuring connections, see “Configuring IP routing connections” on page 2-30.)

The network diagram in Figure 2-4 shows a static route to a subnet specified in the LAN Adrs parameter (10.9.8.10/22) of a Connection profile. With this LAN Adrs parameter setting, the implied static route is defined with the following addresses:

- Dest=10.9.8.10/22

- Gateway=10.9.8.10




Figure 2-4. An IP routing connection serving as a static route

Note: If you do not specify the netmask in the LAN Adrs parameter, the Pipeline inserts a default netmask which assumes the entire far-end network is accessible. Normally, if the far-end router's address includes a netmask, you should include it.

When RIP is turned off in a Connection profile, the Pipeline does not listen to RIP updates across that connection. To route to other networks through that connection, it must rely on a Static Rtes profile. The network diagram in Figure 2-5 shows a remote network that does not have its own Connection profile, but can be reached through an existing Connection profile.




Figure 2-5. When a two-hop static route is required with RIP off

In the example network shown in Figure 2-5, if RIP is off in the Connection profile for site B, the Pipeline must have a Static Rtes profile to site C. A sample profile is shown below:

```
Name=sitec-net
Active=Yes
Dest=10.4.5.6/22
Gateway=10.9.8.10
```



```
Metric=2
Private=Yes
```

Creating a Static Rtes profile

To configure a Static Rtes profile:

- 1 Open the Ethernet > Static Rtes > *any profile*.

- 2 Assign the route a name.

For example:

```
Name=sales-gw
```

- 3 Specify that the route should be added to the routing table.

```
Active=Yes
```

- 4 Specify the destination network.

For example:

```
Dest=10.210.1.30/12
```

The Pipeline must have a Connection profile that specifies this address.

If the address includes a netmask, the remote router is seen as a gateway to that subnet, rather than to a whole remote network. To specify the entire remote network, you would use a network address such as:

```
Dest=10.0.0.0
```

- 5 Specify the address of the router to use for that destination.

For example:

```
Gateway=10.9.8.10
```

This parameter states that the path to the destination subnet is through the IP router at 10.9.8.10.

- 6 Specify a metric for this route.

For example:

```
Metric=1
```

RIP uses distance vector metrics, so the metric is interpreted like a hop count. If the Pipeline has more than one possible route to a destination network, it chooses the one with the lower metric.

- 7 Specify whether this route is private.

For example:

Private=No

This setting specifies that the Pipeline will disclose the existence of the route when queried by RIP or another routing protocol.

- 8 Close and save the profile.

Configuring the default route

If no routes exist for the destination address of a packet, the Pipeline forwards the packet to the default route. Most sites use the default route to specify a local IP router (such as a UNIX host running the route daemon). This helps to offload routing tasks to other devices.

Note: If there is no default route, the Pipeline drops packets for which it has no route. By default, the Pipeline uses the value you entered for the Rem Adr parameter in the Configure profile as the default gateway.

To configure the default route:

- 1 Open the Ethernet > Static Rtes > Default profile.
The name of that profile is always Default, and its destination is always 0.0.0.0 (you cannot change these values).
- 2 Specify that the route should be added to the routing table.
Active=Yes
- 3 Specify the address of the router to use for packets with unknown destinations.
For example:
Gateway=10.9.8.10
- 4 Specify a metric for this route.
For example:
Metric=1
- 5 Specify whether this route is private.
For example:
Private=Yes
This setting specifies that the Pipeline will not disclose the existence of the route when queried by RIP or another routing protocol.
- 6 Close and save the Default profile.

Specifying default routes on a per-user basis

You can specify a default route on a per-user basis by setting the parameter in Ethernet > Connection > *profile* > IP Options > Client Gateway. When the IP address of the user's default route is set, the Pipeline routes IP packets in this way:

- 1 The Pipeline consults its routing table to find a next-hop address.
- 2 If the next hop is the default route for the system (destination 0.0.0.0), the Pipeline uses the per-user default address as a next hop instead of the system-wide default route.

The unit also uses the per-user default if the normal routing logic fails to find a route and there is no system-wide default route.

The Client Gateway IP address applies to routing all packets received on an interface using that profile, regardless of the specific IP source address; therefore, you can set this parameter when the profile belongs to another access router and all hosts behind that router use the default gateway. While all packets arriving on the interface using the given profile are affected, the Pipeline handles packets from other users or from the Ethernet normally. In addition, this feature does not alter the global routing table.

To configure a per-user route in the Pipeline configuration interface, you must set the Client Gateway parameter in the IP Options menu of the Connection profile.

For example:

```
Ip options...
LAN Adrs=nnn.nnn.nnn.nnn/nn
WAN Alias=0.0.0.0
IF Adrs=0.0.0.0/0
Preference=60
Metric=1
DownPreference=120
DownMetric=7
Private=No
RIP=Off
Client Pri DNS=0.0.0.0
Client Sec DNS=0.0.0.0
Client Assign DNS=Yes
>Client Gateway=10.0.0.3
```

Enabling the Pipeline to use dynamic routing

In addition to RIP, the Pipeline can use Internet Control Message Protocol (ICMP) Redirects to acquire routes dynamically. ICMP dynamically determines the best IP route to a destination network or host and uses ICMP redirect packets to transfer packets over a more efficient route. ICMP redirect packets are one of the oldest route discovery methods on the Internet and one of the least secure, due to the possibility of receiving counterfeit ICMP redirects. You can configure the Pipeline to ignore ICMP redirects to promote security.

To ignore ICMP redirects:

- 1 Open the Ethernet > Mod Config menu.
- 2 Make sure that ICMP redirects are not accepted.
`ICMP Redirects=Ignore`
- 3 Close and save the profile.

If you are using RIP-v1

The Internet Engineering Task Force (IETF) voted to move RIP-v1 into the “historic” category so its use is no longer recommended. You can upgrade all routers and hosts to RIP-v2. If you need to maintain RIP-v1, create a separate subnet and place all RIP-v1 routers and hosts on that subnet.

Note: RIP Policy and RIP Summary are relevant only to RIP-v1 and should not be set when interacting with RIP-v2 routers.

If the Pipeline Ethernet interface is on a RIP-v1 subnet:

- 1 Open the Ethernet > Mod Config > Ether Options menu.
- 2 Turn on RIP-v1.
For example:

`RIP=Both-v1`

This setting means that the Pipeline transmits and receives RIP-v1 updates on the local Ethernet. If you do not want the Pipeline to be informed about local routing changes (for example, if all local routing is handled by a default router), you can use the following setting instead:

`RIP=Send-v1`

Or, if you do not want the Pipeline to transmit its WAN connections to the RIP-v1 routers on the local subnet:

`RIP=Recv-v1`

3 Set Ignore Def Rt to Yes.

The default route specifies a static route to another IP router, which is often a local router such as another Pipeline. When the Ignore Def Rt parameter is set to Yes (recommended), RIP updates do not modify the default route in the Pipeline routing table.

4 Close and save the profile.

Configuring RIP-v2 on Ethernet

To turn on RIP-v2 on the local Ethernet:

1 Open the Ethernet > Mod Config > Ether Options menu.

2 Turn on the RIP parameter.

For example:

`RIP=Both-v2`

This setting means that the Pipeline transmits and receives RIP-v2 updates on the local Ethernet. If you do not want the Pipeline to be informed about local routing changes (for example, if all local routing is handled by a default router), you can use the following setting instead:

`RIP=Send-v2`

3 Set Ignore Def Rt to Yes.

The default route specifies a static route to another IP router, which is often a local router such as a Cisco or another Pipeline. When the Ignore Def Rt parameter is set to Yes (recommended), RIP updates will not modify the default route in the Pipeline routing table.

4 Close and save the profile.

Configuring RIP for incoming WAN connections

Many sites turn off RIP on the WAN interface because it can cause very large local routing tables. If RIP is enabled to both send and receive RIP updates over the WAN interface, the Pipeline broadcasts its routing table to the remote network and listens for RIP updates from that network. Gradually, all routers on

both networks implement consistent routing tables (all of which might become quite large).

To configure the Answer profile for RIP and IP routing:

- 1 Open the Ethernet > Answer > PPP Options menu.
- 2 Turn on IP routing.
- 3 Open the Ethernet > Answer > Session Options menu.
- 4 Turn on the RIP parameter.

For example:

```
RIP=Recv-v2
```

This setting means that the Pipeline receives RIP-v2 updates across incoming connections with other IP routers. If you do not want the Pipeline to accept RIP updates on the WAN, use the following settings:

```
RIP=Off
```

- 5 Close and save the Answer profile.

Configuring RIP for a particular connection

You can turn off RIP for a particular connection by configuring it in the Connection profile.

Note: RIP traffic resets the Idle timer and updates are sent every 30 seconds. As such, you should turn off RIP for WAN connections with the Idle (timer) set below 30 seconds, or apply a Call filter for RIP updates on the WAN. If not, the connections will never disconnect.

To configure a Connection profile for RIP and IP routing:

- 1 Open the Ethernet > Connection > *any profile*.
- 2 Turn on IP routing.
- 3 Open the IP Options submenu of the same profile.
- 4 Turn on the RIP parameter.

For example:

RIP=Recv-v2

This setting means that the Pipeline receives RIP-v2 updates from the other IP router.

If the remote router is running RIP-v1 and the local network is running RIP-v2, or if you do not want the Pipeline to send or receive RIP updates on this connection, use the following setting:

RIP=None

- 5 Close and save the Connection profile.

Route preferences

Route preferences provide additional control over which types of routes take precedence over others. For each IP address and netmask pair, the routing table holds one route per protocol, where the protocols are defined as follows:

- Connected routes, such as Ethernet, have a Preference=0.
- Routes learned from ICMP Redirects have a Preference=30.
- Routes placed in the table by SNMP MIB II have a Preference=100.
- Routes learned from RIP have a default Preference=100.

You can modify the default in the Route Preferences submenu of the Ethernet profile.

- A statically configured IP Route or Connection profile has a default Preference=100.

When choosing which routes should be put in the routing table, the router first compares the Preference value, preferring the lower number. If the Preference values are equal, the router then compares the Metric field, using the route with the lower Metric.

If multiple routes exist for a given address and netmask pair, the route with the lower Preference is better. If two routes have the same Preference, then the lower Metric is better. The best route by these criteria is actually used by the router. The others remain latent or “hidden,” and are used in case the best route was removed.

To control route preferences, you can enter a lower (better) preference value using any of the following parameters:

- Ethernet > Connections > *any profile* > IP options > Preference=[]
- Ethernet > Static Rtes > *any profile* > Preference=[]
- Ethernet > Mod Config > Route Pref
Static Preference=100
Rip Preference=100

Viewing the routing table

The `iproute show terminal-server` command includes information relevant to multiple IP routing protocols. To view the IP routing table, invoke the terminal server interface and at the prompt, enter:

```
iproute show
```

The output looks similar to the following table:

Destination	Gateway	IF	Flg	Pref	Met	Use	Age
0.0.0.0/0	10.0.0.100	wan0	SG	1	1	0	20887
10.207.76.0/24	10.207.76.1	wanidle0	SG	100	7	0	20887
10.207.76.1/32	10.207.76.1	wanidle0	S	100	7	2	20887
10.207.77.0/24	10.207.76.1	wanidle0	SG	100	8	0	20887
127.0.0.1/32	-	lo0	CP	0	0	0	20887
10.0.0.0/24	10.0.0.100	wan0	SG	100	1	21387	20887
10.0.0.100/32	10.0.0.100	wan0	S	100	1	153	20887
10.1.2.0/24	-	ie0	C	0	0	19775	20887
10.1.2.1/32	-	ie0	CP	0	0	389	20887
255.255.255.255/32	-	ie0	CP	0	0	0	20887

The column headings shown here are described in “Fields in the routing table” on page 2-27. The routes in this table are explained as follows:

Destination	Gateway	IF	Flg	Pref	Met	Use	Age
0.0.0.0/0	10.0.0.100	wan0	SG	1	1	0	20887

This is the default route, pointing through the active Connection profile. The Static Rtes profile for the default route specifies a Preference of 1, so this route is preferred over dynamically learned routes.

Destination	Gateway	IF	Flg	Pref	Met	Use	Age
10.207.76.0/24	10.207.76.1	wanidle0	SG	100	7	0	20887
10.207.76.1/32	10.207.76.1	wanidle0	S	100	7	2	20887

Configuring IP Routing

Managing the routing table

These routes are specified in a Connection profile. Note that there are two routes—a direct route to the gateway itself and a route to the larger network.

Destination	Gateway	IF	Flg	Pref	Met	Use	Age
10.207.77.0/24	10.207.76.1	wanidle0	SG	100	8	0	20887

This is a static route that points through an inactive gateway.

Destination	Gateway	IF	Flg	Pref	Met	Use	Age
127.0.0.1/32	–	lo0	CP	0	0	0	20887

This is the loopback route, which says that packets sent to this special address will be handled internally. The C flag indicates a Connected route, while the P flag indicates that the router will not advertise this route.

Destination	Gateway	IF	Flg	Pref	Met	Use	Age
10.0.0.0/24	10.0.0.100	wan0	SG	100	1	21387	20887
10.0.0.100/32	10.0.0.100	wan0	S	100	1	153	20887

These routes are created by a Connection profile that is currently active. These are similar to the 10.207.76.0 routes shown above, but these routes live on an active interface.

Destination	Gateway	IF	Flg	Pref	Met	Use	Age
10.1.2.0/24	–	ie0	C	0	0	19775	20887

This route describes the connection to the Ethernet interface. It is directly connected, with a Preference and Metric of zero.

Destination	Gateway	IF	Flg	Pref	Met	Use	Age
10.1.2.1/32	–	ie0	CP	0	0	389	20887

This is another loopback route, a host route with the local Ethernet address. It is private, so it will not be advertised.

Destination	Gateway	IF	Flg	Pref	Met	Use	Age
255.255.255.255/32	–	ie0	CP	0	0	0	20887

This is a private route to the broadcast address. It is used in cases where the router needs to broadcast a packet but is otherwise unconfigured. The route is typically used when trying to locate a server on a client machine to handle challenges for a token security card.

Fields in the routing table

The columns in the routing table display the following information:

- Destination

The Destination column indicates the target address of a route. To send a packet to this address, the Pipeline will use this route. Note that the router will use the most specific route (having the largest netmask) that matches a given destination.

- Gateway

The Gateway column specifies the address of the next hop router that can forward packets to the given destination. Direct routes (without a gateway) do not have a gateway address.

- IF

The Interface column shows the name of the interface through which a packet addressed to this destination will be sent.

- bh0 is the black-hole interface. It has an IP address of 127.0.0.3. Packets routed to this interface are discarded silently.
- ie0 is the Ethernet interface.
- lo0 is the loopback interface.
- local. Routes pointing to local machines are labeled local. These include the following routes, with a single w.x.y.z route for each local IP address:

127.0.0.1/32	-	local	CP	0	0	0	59593
224.0.0.1/32	-	local	CP	0	0	0	59593
224.0.0.2/32	-	local	CP	0	0	0	59593
w.x.y.z/32	-	local	CP	0	0	0	59593

- mcast. Routes to 224.0.0.1 and 224.0.0.2 represent the multicast addresses for all systems on the local subnet and all routers on the local subnet, respectively, and are never forwarded. All multicast addresses (except for addresses 224.0.0.1/32 and 224.0.0.2/32) point to the mcast interface.

224.0.0.0/4	-	mcast	CP	0	0	0	59593
-------------	---	-------	----	---	---	---	-------

- rj0 is the reject interface. It has an IP address of 127.0.0.2. Packets routed to this interface are sent back to the source address with the ICMP “host unreachable” message.
- wann specifies one of the active WAN interfaces.
- wanidle0 is the inactive interface (the special interface where all routes point when their WAN connections are down).
- Flg
The Flg column can contain the following flag values:
 - C=Connected (A directly connected route. For example, the Ethernet.)
 - I=ICMP (ICMP Redirect dynamic route.)
 - N=NetMgt (Placed in the table via SNMP MIB II.)
 - R (A RIP dynamic route.)
 - S=Static (A locally configured Static Rtes profile or Connection profile route.)
 - ?=Unknown (Indicates an error.)
 - G=Gateway (A gateway is required in order to reach this route.)
 - P=Private (This route will not be advertised via RIP.)
 - T=Temporary (This route will be destroyed when its interface goes down.)
 - *=Hidden (A hidden route means that there is a better route in the table, so this route is hidden “behind” the better route. If the better route goes down, then this route might be used.)
- Pref
The Preference column contains the preference value of the route. Note that all routes that come from RIP will have a preference value of 100, while the preference value of each individual static route may be set independently. (To set a route independently, see “Route preferences” on page 2-24.)
- Metric
The Metric column shows the RIP-style metric for the route, with a valid range of zero to 16.
- Use

This is a count of the number of times the route has been referenced since it was created. (Many of the references are internal, so this is not a count of the number of packets sent using this route.)

Unused routes are indicated by a 0 in the Use column.

- **Age**

This is the age of the route in seconds. It is used for troubleshooting, to determine when routes are changing rapidly (referred to as “flapping”).

Removing down routes to a host

The Pipeline advertises addresses associated with Connection profiles as routes to which it can connect. By default, it advertises these addresses even when a link is down, because they are necessary for the on-demand connections that the Pipeline establishes.

For a nailed connection, it is assumed that the connection is always up. If it is not, the routes to that connection are not necessary until the connection comes back up. For example:

Pipeline 1 and Pipeline 2 are on the same local LAN.

- Pipeline 1 has a nailed connection to a remote site.

The remote address has a metric of 4.

- Pipeline 2 is a backup connection.

It has a remote address with a metric of 7.

Traffic goes through Pipeline 1 because of the lower metric. If its connection goes down, its route to the remote network is still advertised by default. Therefore, the connection specified by Pipeline 2 never comes up.

To remove the route of a down, nailed connection, set the Temporary parameter in Ethernet > Connection > *profile of down connection* > IP options submenu to Yes. When the Temporary parameter is set to Yes, a route to a nailed connection is removed from the routing table when the link is down, including all routes dynamically learned on this connection, and discontinues advertising the route. The routes are advertised and reappear in the routing table only when you re-establish the connection.

Identifying Temporary routes in the routing table

The “T” flag appears in the IP routing display to indicate temporary routes. In this example, the Show IP Routes command displays two temporary routes:

```
ascend% show ip routes
```

Destination	Gateway	IF	Flg	Pref	Met	Use	Age
192.168.252.0/30	192.168.252.1	wan10	rGT	60	7	0	7
192.168.252.1/32	192.168.252.1	wan10	rT	60	7	1	7

Configuring IP routing connections

Note: If you configure a routing configuration to a second destination, be sure to specify routing information for both sides. Specify the remote network information in the Connection profile for that network. Network information for the local Ethernet is configured in the Ethernet > Mod Config profile.

This section describes how to configure IP routing connections. It describes typical host software requirements and includes the following example configurations:

- Example host connection with static address
- Example router connection
- Example router connection on a subnet

Note: The most common cause of trouble in initially establishing an IP connection is incorrect configuration of the IP address or subnet specification for the remote host or calling device.

Checking remote host requirements

IP hosts, such as UNIX systems, Windows or OS/2 PCs, or Macintosh systems, must have appropriately configured TCP/IP software. A remote host calling into the local IP network must also have PPP software.

- UNIX
UNIX systems typically include a TCP/IP stack, DNS software, and other software, files, and utilities used for Internet communication. UNIX network

administration documentation describes how to configure these programs and files.

- **PC-compatibles**
PCs running Windows or OS/2 need the TCP/IP networking software or “stack.” The stack is included with Windows 95, but the user might have to purchase and install it separately if the computer has a previous version of Windows or OS/2.
- **Macintosh**
Macintosh computers need MacTCP or Open Transport software for TCP/IP connectivity. MacTCP is included with all Apple system software including and after Version 7.1. (You can see if the software is present by looking in the Control Panel folder for MacTCP or MacTCP Admin.)

For any platform, the TCP/IP software must be configured with the host’s IP address and subnet mask. If a DNS server is supported on your local network, you should also configure the host software with the DNS server’s address. Typically, the host software is configured so the Pipeline is the default gateway. (Refer to the *Start Here Guide* for TCP/IP configuration examples.)

Also see a discussion about how the Pipeline translates an acquired IP address from a Network Address Server (NAS) and manages traffic between hosts on the local network and the wide area network in “Network Address Translation (NAT) for a LAN” on page 3-22.

Example host connection with static address

A host route connection enables the dial-in host to keep its own IP address when logging into the Pipeline IP network. For example, if a PC user telecommutes to one IP network and uses an ISP on another IP network, one of those connections can assign an IP address and the other can configure a host route to the PC. The

following shows how to configure a host route. (For details on the /32 netmask, see “Subnet mask notation” on page 2-4.)




Figure 2-6. A dial-in user requiring a static IP address (a host route)

In this example, the PC on the right in Figure 2-6 is running PPP software, a TCP/IP stack, and has an ISDN modem card. The PPP software includes settings like these:

```
Username=Simon
Accept Assigned IP=N/A (or No)
IP address=10.8.9.10
Netmask=255.255.255.255
Default Gateway=N/A (or None)
Name Server=10.7.7.1
Domain suffix=abc.com
VAN Jacobsen compression ON
```

To configure the Pipeline to accept dial-in connections from this host:

- 1 Open the Ethernet > Answer > PPP options menu.
- 2 Make sure that IP Routing is enabled.
For example:
Route IP=Yes
- 3 Close and save the profile.
- 4 Open Ethernet > Connection > profile for Simon.
- 5 Set these parameters:

```
Station=Simon
Active=Yes
Encaps=PPP
Route IP=Yes
```

```
Encaps options...
  Send Auth=CHAP
  Recv PW=*SECURE*

IP options...
  LAN Adrs=10.8.9.10/32
  RIP=Off
```

- 6 Close and save the profile.

Example router connection

In the following example, the Pipeline is connected to a corporate IP network, and needs a switched connection to another company that has its own IP configuration. Figure 2-7 shows an example network diagram.




Figure 2-7. A router-to-router IP connection

This example assumes that the Ethernet > Answer profile and Ethernet > Mod Config > Ether options have been set up properly to enable IP routing.

To set up a Connection profile on the Pipeline at site A to link to site B, do the following:

- 1 Open Ethernet > Connection > profile for site B.
- 2 Set these parameters:

```
Station=PipelineB
Active=Yes
Encaps=MPP
Route IP=Yes

Encaps options...
  Send Auth=CHAP
```


Configuring IP Routing

Configuring IP routing connections

```
Recv PW=*SECURE*
Send PW=*SECURE*

IP options...
LAN Adrs=10.9.8.7/22
RIP=Send-v2
```

- 3 Close and save the profile.

To configure the Pipeline at site B link to the one at site A, do the following:

- 1 Open Ethernet > Connection > profile for site A.
- 2 Set these parameters:

```
Station=PipelineA
Active=Yes
Encaps=MPP
Route IP=Yes

Encaps options...
Send Auth=CHAP
Recv PW=*SECURE*
Send PW=*SECURE*

IP options...
LAN Adrs=10.2.3.1/22
RIP=Recv-v2
```

- 3 Close and save the profile.

Example router connection on a subnet

In the following example network, the Pipeline is used to connect telecommuters with their own Ethernet networks to the corporate backbone. The Pipeline is on a subnet, and assigns subnet addresses to the telecommuters' networks.




Figure 2-8. A connection between local and remote subnets

This example assumes that the Ethernet > Answer profile and Ethernet > Mod Config > Ether options in both devices have been set up properly to enable IP routing.

Because the Pipeline specifies a netmask as part of its own IP address, it must use other routers to reach any IP addresses outside that subnet. To forward packets to other parts of the corporate network, the Pipeline must either have a default route configuration to a router in its own subnet, or must enable RIP on Ethernet.

To configure the Pipeline at site A with an IP routing connection to site B:

- 1 Open the Ethernet > Connection > profile for site B.
- 2 Set these parameters:

```
Station=PipelineB
Active=Yes
Encaps=MPP
Route IP=Yes

Encaps options...
  Send Auth=CHAP
  Recv PW=*SECURE*
  Send PW=*SECURE*
```

Configuring IP Routing

Configuring IP routing connections

```
IP options...
```

```
LAN Adrs=10.7.8.200/24
```

```
RIP=Off
```

- 3 Close and save the profile.
- 4 Open the Ethernet > Static Rtes > Default profile.
- 5 Set these parameters:

Active=Yes
Gateway=10.4.4.133/24
Metric=1
Preference=100
Private=Yes
- 6 Close and save the profile.

On the site B router:

- 7 Open Ethernet > Connection > profile for site A.
- 8 Set these parameters:

Station=MAXA
Active=Yes
Encaps=MPP
Route IP=Yes
Encaps options...
 Send Auth=CHAP
 Recv PW=*SECURE*
 Send PW=*SECURE*
IP options...
 LAN Adrs=10.4.5.1/24
 RIP=Off
- 9 Close and save the profile.
- 10 Open the Ethernet > Static Rtes > Default profile on the site B Pipeline.
- 11 Set these parameters:

Active=Yes
Gateway=10.4.5.1/24
Metric=1
Preference=100
Private=Yes

12 Close and save the profile.

Ascend Tunnel Management Protocol (ATMP)

Virtual private networks can include the Pipeline as a Home Agent ATMP end point in implementations where the Pipeline operates in router mode.

Using a Pipeline in a virtual private network

Virtual private networks provide low-cost remote access to private LANs via the Internet. The tunnel to the private corporate network might be from an ISP, enabling mobile nodes to dial into a corporate network, or between two corporate networks that access one another through a low-cost Internet connection.

Ascend Tunnel Management Protocol (ATMP) uses a UDP/IP session between two units to build a tunnel for encapsulated packets. It puts the packets in standard Generic Routing Encapsulation (GRE), as described in RFC 1701. In effect, the tunnel collapses the Internet cloud and provides what looks like direct access to a home network. The packets must be routed (IPX or IP).

Foreign and home agents

ATMP tunnels work between two Ascend units. One of the units acts as a foreign agent (typically a local ISP) and one as a home agent (which can access the home network). A mobile node dials into the foreign agent, which establishes a cross-Internet IP session with the home agent. The foreign agent then requests an ATMP tunnel on top of the IP session. The foreign agent must use RADIUS to authenticate mobile nodes dial-ins.

The home agent is the terminating part of the tunnel, where most of the ATMP intelligence resides. This agent must be able to communicate with the home network (the destination network for mobile nodes) through a direct connection, another router, or across a nailed connection.

The home agent may communicate with the home network through a direct connection, another router, or across a nailed connection. When it relies on

Configuring IP Routing

Ascend Tunnel Management Protocol (ATMP)

packet routing to reach the home network, it operates in router mode. It is in gateway mode when it has a nailed connection to the home network.

A home agent can be an Ascend MAX or a Pipeline 50 or 130. When a Pipeline is used as the home agent end point, only routing is supported.

Configuring a home agent in router mode

With the ATMP tunnel established between the home agent and foreign agent, the home agent receives IP packets through the tunnel, removes the GRE encapsulation, and passes the packets to its bridge/router software. It also adds to its routing table, a host route to the mobile node.

Following are the parameters for configuring a home agent in router mode. The IPX routing parameters in the Ethernet profile are required only if the Pipeline is routing IPX.

```
Ethernet
Mod Config
  IPX Routing=Yes
  Ether options...
    IP Adrs=10.1.2.3/24
    IPX Frame=802.2
    IPX Enet #=00000000
  ATMP options...
    Password=private
    UDP Port=5150
```

Password is the password used to authenticate the ATMP tunnel itself. It must match the password specified by the Ascend-Home-Agent-Password attribute of the mobile nodes' RADIUS profiles. (All mobile nodes use the same password for that attribute.)

ATMP uses UDP port 5150 for ATMP messages between the foreign and home agents. If you specify a different UDP port number, make sure that the entire ATMP configuration agrees.

Following are the parameters for the IP routing connection to the foreign agent, which is authenticated and established in the usual way:

Ethernet

Connections

Station=foreign-agent

Active=Yes

Encaps=MPP

Dial #=555-1213

Route IP=Yes

Encaps options...

Send Auth=CHAP

Recv PW=foreign-pw

Send PW=home-pw

IP options...

LAN Adrs=10.65.212.226/24

IP Address Management

This chapter includes the following topics:

Connecting to a local IP network	3-1
BOOTP Relay	3-9
DHCP services	3-10
Dial-in user DNS server assignments	3-15
Local DNS host address table	3-17
Network Address Translation (NAT) for a LAN	3-22

Connecting to a local IP network

To connect the Pipeline to your local IP network, you need to assign the Pipeline Ethernet interface an IP address. In addition, you might want to perform one or more of the following tasks:

- Enable proxy ARP to let the Pipeline respond to ARP requests for remote nodes.
- Configure DNS or WINS information to enable users to Telnet in using host names.
- Configure the Pipeline to generate UDP checksums.
- Update other IP routers on the backbone.

The list below shows the relevant configuration parameters:

- Ethernet > Mod Config > Ether Options

```
IP Adrs=10.2.3.1/24
2nd Adrs=10.128.8.55/24
RIP=Both-v2
RIP2 Use Multicast=Yes
Ignore Def Rt=Yes
Proxy Mode=Off
UDP Cksum=Yes
TCP Timeout=100
```

- Ethernet > Mod Config > DNS

```
>Domain Name=abc.com
Sec Domain Name=Yes
Allow As Client DNS=Yes
List Attempt=Yes
List Size=6
Client Pri DNS=0.0.0.0
Client Sec DNS=0.0.0.0
```

If the DNS system is set up to return lists of host addresses in response to a query, the List Attempt parameter enables a user to attempt a login to one entry in the DNS list of hosts, and if that connection fails, to try the next entry, and so on. This helps to avoid tearing down physical links when a host is unavailable, which is especially important for immediate services such as immediate Telnet or Rlogin.

The List Size parameter specifies a number of addresses that will be listed. The maximum number is 35. Also see “User-definable TCP connection retry timeout” on page 3-21 to use the TCP Timeout parameter to attempt subsequent DNS servers, as needed.

- Ethernet > Static Rtes > *any profile*

```
Name=xyz.com
Active=Yes
Dest=198.2.3.0/24
Gateway=198.2.3.4
Metric=2
Preference=100
Private=No
```

For details on each parameter, see the *Reference Guide*, and for information about using RIP on Ethernet, see “Enabling the Pipeline to use dynamic routing” on page 2-21.

Assigning the Ethernet interface IP address

The Pipeline Ethernet interface must have a unique IP address that is consistent with the addresses of other hosts and routers on the same network.

To assign the Pipeline an IP address on the Ethernet:

- 1 Open the Ethernet > Mod Config > Ether Options menu.
- 2 Enter the IP address for the Ethernet interface in IP Adrs.

For example:

IP Adrs=10.2.3.1

- 3 Close and save the profile.

After you have configured the IP address, you can Ping the Pipeline from a host to verify that it is up and running on the network. (How to use the Ping command is described in “Using Ping to verify the address” on page 3-6.)

Creating a subnet for the Pipeline

On a large corporate backbone, administrators often configure subnets to increase the network address space, segment a complex network, and control routing in the local environment. For example, suppose the main backbone IP network is 10.0.0.0, and supports a router at 10.0.0.17.




Figure 3-1. Creating a subnet for the Pipeline

You can place the Pipeline on a subnet of that network by entering a subnet mask in its IP address specification, for example:

- 1 Open the Ethernet > Mod Config > Ether Options menu.
- 2 Enter the IP address for the Ethernet interface in the IP Adrs field.

For example:

IP Adrs=10.2.3.1/24

- 3 Close and save the profile.

With this subnet address, the Pipeline requires a static route to the backbone router on the main network. Otherwise, it can only reach the subnets to which it is directly connected.

To create the static route and make the backbone router the default route:

- 1 Open the Ethernet > Static Rtes > Default profile.
- 2 Specify the IP address of a backbone router in the Gateway field.

For example:

Gateway=10.0.0.17

- 3 Leave the other parameters at their default values.

For example:

Active=Yes

Dest=0.0.0.0/0

Metric=1

Private=Yes

- 4 Close and save the profile.

Assigning two addresses: Dual IP

The Pipeline can assign two separate IP addresses to a single physical Ethernet port and route between them—a feature often referred to as “dual IP.” The two addresses provide logical interfaces to two networks or subnets on the same backbone.

Usually devices connected to the same physical wire belong to the same IP network. With dual IP, one wire can support two IP networks. Devices on the wire are assigned to one network or the other. The devices route information to each other through the Pipeline.

Dual IP is also used to distribute the load of routing traffic to a large subnet by assigning IP addresses on that subnet to two or more routers on the backbone. With a direct connection to the subnet as well as to the backbone network, each of the routers routes packets to devices on the subnet and includes the route in their routing table updates.

Dual IP also allows you to make a smooth transition when changing IP addresses. That is, a second IP address can act as a place holder while IP addresses are changed on other network equipment.

Figure 3-2 shows two routers configured with a second address on the same subnet.




Figure 3-2. Dual IP and shared subnet routing

To assign two addresses to the Pipeline Ethernet interface:

- 1 Open the Ethernet > Mod Config > Ether Options menu.
- 2 Enter the IP address for the Ethernet interface in the IP Adrs field.

For example:

IP Adrs=10.2.3.1/24

- 3 Type the second IP address in the 2nd Adrs field.

For example:

IP Adrs=10.128.8.55/24

After you have configured the IP addresses, you can Ping the addresses from another IP host on each of the IP subnets to verify that both logical interfaces are accessible.

Note: For other routers to recognize the Pipeline on either of its two networks, you must either turn on RIP on the Ethernet interface or configure static routes in those routers.

- 4 Close and save the profile.

Using Ping to verify the address

The Ping command sends an Internet Control Message Protocol (ICMP) mandatory echo request datagram, which asks the remote station “Are you there?” If the echo request reaches the remote station, the station sends back an ICMP echo response datagram, which tells the sender, “Yes, I am alive.” This exchange verifies that the transmission path is open between the Pipeline and another station.

To verify that the Pipeline is up on the local network, invoke the terminal server interface and enter this command:

```
ping <host-name>
```

For example:

```
ping 10.1.2.3
```

You can terminate the Ping exchange at any time by pressing Ctrl-C. (For more information about verifying that a device is on the network, see Chapter 8, “Pipeline System Administration.”)

Enabling proxy mode in the Pipeline

When a dial-in host has an IP address on the same network as the Pipeline, only the Pipeline keeps track that packets addressed to the host must be routed across the WAN. To other local routers and hosts, the address appears to be on the local network. Therefore, they might broadcast Address Resolution Protocol (ARP) requests on the local network expecting the apparently local host to respond with its physical address. Because the host is not really local, it cannot receive the requests. But if the Pipeline is in Proxy Mode, serving as a proxy for the remote host, it responds with its own physical address.

To enable the Pipeline to respond to ARP requests for remote devices that have local IP addresses:

- 1 Open the Ethernet > Mod Config > Ether Options menu.
- 2 Turn on Proxy Mode.

If the IP addresses are assigned dynamically, use this setting:

```
Proxy Mode=Active
```

If the IP addresses are assigned statically, use this setting instead:

`Proxy Mode=Always`

- 3 Close and save the profile.

Enabling DNS on the Pipeline

If the local network supports Domain Name System (DNS) servers, you can configure the local domain name and the IP addresses of those servers in the Ethernet profile.

If the Pipeline is configured for DNS, users can execute TCP/IP commands such as Telnet and Ping from the Pipeline terminal server interface with host names instead of IP addresses. In addition, the List Attempt parameter helps avoid tearing down physical links by enabling the user to try one entry in the DNS list of hosts, and if that connection fails, to try the next entry, and so on.

To configure the Pipeline for DNS:

- 1 Open the Ethernet > Mod Config > DNS menu.
- 2 Enter your domain name.
For example:
`Domain Name=eng.abc.com`
- 3 Specify the IP address of the primary and secondary DNS servers.
For example:
`Pri DNS=10.2.3.56`
`Sec DNS=10.2.3.107`
- 4 If your site supports multiple addresses for a DNS host name, turn on List Attempt.
`List Attempt=Yes`
- 5 Close and save the profile.

Generating UDP checksums

User Datagram Protocol (UDP) supports the optional use of a checksum field for checking the integrity of both the UDP header and data. The Pipeline always checks the UDP checksum field of each UDP packet it receives, and generates

Ethernet and PPP checksums for the appropriate packets. However, it does not generate UDP checksums unless you set the UDP Cksum parameter.

You should turn on UDP checksums if data integrity is of the highest concern for your environment and you need redundant checks. UDP checksums are also appropriate if your UDP-based servers are located on the remote side of a WAN link that is prone to errors.

Currently the Pipeline uses UDP when generating queries and responses for the following protocols:

- SYSLOG
- DNS
- ECHOSERV
- RIP
- TFTP

To configure the Pipeline to generate checksums for these packets:

- 1 Open the Ethernet > Mod Config menu.
- 2 Turn on UDP checksums.
UDP Cksum=Yes
- 3 Close and save the profile.

Updating other routers on the backbone

If you want to update the routing tables of other local routers whenever the Pipeline brings up a remote connection, configure the Pipeline to send RIP updates over the Ethernet interface. The Pipeline then broadcasts RIP packets containing information about each route change. RIP updates are sent every 30 seconds, so within a minute or so, all routers on the local network are informed about the new route. You can also configure the Pipeline to receive RIP updates on Ethernet, or to both send and receive the updates. (For instructions, see “Configuring RIP-v2 on Ethernet” on page 2-22.)

BOOTP Relay

The Bootstrap Protocol (BOOTP) defines how a computer on a TCP/IP network can get its Internet Protocol (IP) address and other information it needs to start up from another computer. The computer that requests startup information is called the BOOTP client, and the computer that supplies the startup information is called the BOOTP server. A request for startup information is called a BOOTP request, and the BOOTP server's response is called a BOOTP reply.

When the BOOTP client and BOOTP server are not on the same local-area network, the BOOTP request must be relayed from one network to another. This task, known as BOOTP relay, can be performed by a Pipeline.

A device that relays BOOTP requests to another network is known as a BOOTP relay agent. In addition to delivering BOOTP requests to servers, a BOOTP relay agent is responsible for delivering BOOTP replies to clients. In most cases, the agent is a router that connects the networks, such as a Pipeline.

By default, a Pipeline does not relay BOOTP requests to other networks. To enable the BOOTP relay feature for BOOTP clients connected to your Pipeline, follow these steps:

- 1 Obtain the IP address of up to two BOOTP servers to be used.
- 2 Open the Ethernet > Mod Config:

```
20-A00 Mod Config
BOOTP Relay...
>BOOTP Relay Enable=No
Server=0.0.0.0
Server=0.0.0.0
```
- 3 Select BOOTP Relay Enable and set it to Yes.
- 4 Select Server and press Enter to open a text box. In the text box, enter the IP address of the BOOTP server. Press Enter to close the text box.
- 5 If there is another BOOTP server available, select the second menu item named Server and enter its IP address.

You are not required to specify a second BOOTP server.

Note: If you specify two BOOTP servers, the Pipeline that relays the BOOTP request determines when each server is used. The order of the

BOOTP servers in the BOOTP Relay menu does not necessarily determine which server is tried first.

Note: Previously, the Pipeline could not enable both BOOTP relay and DHCP spoofing at the same time because the two functions attempted to respond to the same packets in different ways. Now, if both features are enabled and no WAN links are active, the Pipeline performs DHCP spoofing. As soon as the dialed link is established, the Pipeline stops spoofing and acts as a BOOTP relay agent.

DHCP services

A Pipeline can perform a number of Dynamic Host Configuration Protocol (DHCP) services, including:

- DHCP Server functions, responding to DHCP requests for up to 43 clients at any given time. DHCP server responses provide an IP address and subnet mask. Two address pools of up to 20 IP addresses each can be defined. Additionally, up to three hosts, identified by their MAC (Ethernet) addresses, can have an IP address reserved for their exclusive use.
- Managing Plug and Play requests for TCP/IP configuration settings from computers using Microsoft Windows 95 or Windows NT.
- DHCP Spoofing responses, supplying a temporary IP address for a single host. The IP address supplied is always one greater than that of the Pipeline. The IP address is good for only 60 seconds—just long enough to allow a security-card user to acquire the current password from an ACE or SAFEWORLD server and bring up an authenticated dial-up session. Once the dial-up session is established, an official IP address can be retrieved from a remote DHCP or BOOTP server.

How IP addresses are assigned

When a Pipeline is configured to be a DHCP server and it receives a DHCP client request, it assigns an IP address in one of the following ways:

- When the plug-and-play option is enabled (DHCP PNP Enabled=Yes), the Pipeline takes its own IP address, increments it by one, and returns it in the BOOTP reply message along with IP addresses for the Default Gateway and Domain Name Server. Plug-and-play works with Microsoft Windows 95

(and potentially other IP stacks) to assign an IP address and other wide-area networking settings to a requesting device automatically. With plug-and-play you can use the Pipeline to respond to distant networks without having to configure an IP address first.

- If there is an IP address that is reserved for the host, the Pipeline assigns the reserved address.
- If the host is renewing the address it currently has, the Pipeline assigns the host the same address.

When a host gets a dynamically assigned IP address from one of the address pools, it periodically renews the lease on the address until it has finished using it, as defined by the DHCP protocol. If the host renews the address before its lease expires, the Pipeline always provides the same address.

- If the host is making a new request and there is no IP address reserved for the host, the Pipeline assigns the next available address from its address pools. Up to two 20-address pools of contiguous IP addresses are drawn from. Addresses are assigned using the first available address from the first pool or, if there are no available addresses in that pool and there is a second pool, the first available address in the second pool.

Configuring DHCP services

To configure a DHCP service, open Ethernet > Mod Config > DHCP Spoofing.

Set each parameter according to the function it provides, as described in the following list.

Note: Although the name of this menu is DHCP Spoofing, it contains parameters for all DHCP services, including DHCP Spoofing, DHCP Server, and Plug and Play.

```
20-A00 Mod Config
DHCP Spoofing...
  DHCP Spoofing=Yes
  DHCP PNP Enabled=Yes
  Renewal Time=10
  Become Def. Router=No
  Dial If link down=No
  Always Spoof=Yes
  Validate IP=Yes
```

```
Maximum no reply wait=5
IP group 1=181.100.100.100/16
Group 1 count=1
IP group 2=0.0.0.0/0
Group 2 count=0
Host 1 IP=181.100.100.120
Host 1 Enet=0080c75Be95e
Host 2 IP=0.0.0.0/0
Host 2 Enet=000000000000
Host 3 IP=0.0.0.0/0
Host 3 Enet=000000000000
```

- 1 Set the DHCP Spoofing parameter to Yes to enable any DHCP service. This parameter, which was included in earlier versions of the Ascend software, now has a different meaning. It must be Yes for any DHCP service to be enabled. If it is set to No, other settings in this menu are ignored.
- 2 Set the DHCP PNP Enabled parameter to Yes to enable Plug and Plug. Setting this parameter to Yes and DHCP Spoofing set to Yes is all that is required to enable Plug and Play support.
- 3 Renewal Time specifies how long a DHCP IP address lives before it needs to be renewed. It applies to DHCP spoofed addresses and DHCP server replies. If the host renews the address before it expires, the Pipeline provides the same address. Plug and Play addresses always expire in 60 seconds.
- 4 Become Default Router is an option you can set to advertise the address of your Pipeline as the default router for all DHCP request packets.
- 5 Dial If Link Down is used with DHCP spoofing in conjunction with BOOTP Relay. This parameter applies when both DHCP spoofing and BOOTP relay are enabled. If no wide area network links are active, the Pipeline performs DHCP spoofing. When set to Yes, as soon as the dialed link is established, the Pipeline stops DHCP spoofing and acts as a BOOTP relay agent.
- 6 Set Always Spoof as follows:
 - **Yes enables the DHCP server.** A DHCP server always supplies an IP address for every request, until all IP addresses are exhausted.
 - **No enables DHCP spoofing.** DHCP spoofing only supplies an IP address for a single host on the network. It does not respond to all requests.

If both DHCP Spoofing and Always Spoof are Yes, the DHCP server feature is enabled. If DHCP Spoofing is Yes and Always Spoof is No, DHCP spoofing is enabled and works as it did in earlier releases when the value of Always Spoof was Yes.

- 7** Set Validate IP to Yes to check if a spoofed address that is about to be assigned is already in use, and if it is, automatically assign another address.
- 8** Set Maximum No-Reply Wait only if you are validating IP addresses. To validate the IP address, DHCP sends an ICMP echo (ping) to check if the address is in use. The maximum time it waits for a reply is determined by this setting. The default is 10 seconds.
- 9** To assign IP addresses dynamically, set the IP Group 1 parameter to the first address for the IP address pool.
- 10** Set the Group 1 Count parameter to the number of addresses in the pool. The pool can contain up to 20 addresses.
- 11** To define an additional address pool for dynamic address assignment, set the IP Group 2 parameter to the first address for the second IP address pool.
- 12** Set the Group 2 Count parameter to the number of addresses in the pool. The second pool, which can also contain up to 20 addresses, is used only if there are no addresses available in the first pool.
- 13** To reserve an IP address for a particular host, set the Host 1 IP parameter to the IP address for the host.
- 14** Set the Host 1 Enet parameter to the MAC (Ethernet) address of the host. The MAC address is normally the Ethernet address of the network interface card that the host uses to connect to the local-area network. The DHCP server assigns this host the IP address you specify whenever it gets a DHCP request for an IP address from the host with that MAC address.
- 15** To reserve an IP address for another host, set the Host 2 IP parameter to the IP address for the host.
- 16** Set the Host 2 Enet parameter to the MAC (Ethernet) address of the host.
- 17** To reserve an IP address for another host, set the Host 3 IP parameter to the IP address for the host.
- 18** Set the Host 3 Enet parameter to the MAC (Ethernet) address of the host.

Setting up a DHCP server

To set up a DHCP server, these parameters are required to be set:

```
DHCP Spoofing...  
DHCP Spoofing=Yes  
Always Spoof=Yes  
IP group 1=nnn.nnn.nnn.nnn/nn  
Group 1 count=n
```

Additionally, you might set these parameters:

```
Renewal Time=nn  
IP group 2=0.0.0.0/0  
Group 2 count=0  
Host 1 IP=nnn.nnn.nnn.nnn/nn  
Host 1 Enet=0080c75Be95e  
Host 2 IP=0.0.0.0/0  
Host 2 Enet=000000000000  
Host 3 IP=0.0.0.0/0  
Host 3 Enet=000000000000
```

Setting up Plug and Play support

To set up Plug and Play, you must set these parameters:

```
DHCP Spoofing...  
DHCP Spoofing=Yes  
DHCP PNP Enabled=Yes
```

Setting up DHCP spoofing

To set up DHCP spoofing, you must set these parameters:

```
DHCP Spoofing...  
DHCP Spoofing=Yes  
Always Spoof=No
```

Additionally, you might set these parameters:

```
Renewal Time=nn  
Become Def. Router=Yes|No
```

```
Dial If Link Down=Yes|No
Validate IP=Yes
Maximum no reply wait=n
```

Dial-in user DNS server assignments

IP addresses for Domain Name System (DNS) servers can be set for users who dial into the Pipeline via PPP. DNS information is supplied on the basis of these rules:

- First, if Client PRI DNS and Client Sec DNS parameters are specified at the profile level, these parameters are passed to the user.
- Then, if the DNS information is defined in the Ethernet profile, the Pipeline passes these parameters to the user.
- If no client DNS information is defined either at the Connection or Ethernet profile level, and the parameter 'Allow As Client DNS' is set to Yes, the Pipeline passes the primary and secondary (PRI and SEC) DNS information defined for the Pipeline. You can prevent the default DNS information of the Pipeline from being passed to a user when all other IPCP DNS negotiation fails by setting 'Allow As Client DNS' to No.

Configuring DNS servers in the Ethernet profile

To configure user-level DNS servers in the Ethernet profile:

- 1 Open the Ethernet > Mod Config > DNS menu.

For example:

```
30-100 Mod Config
DNS...
Domain Name=
Pri DNS=111.111.111.11
Sec DNS=0.0.0.0
Allow as Client DNS=Yes
List attempt=Yes
List Size=6
Client Pri DNS=101.10.10.1
Client Sec DNS=101.10.10.2
```

```
Enable Local DNS Table=Yes  
Loc. DNS Tab Auto Update=Yes
```

- 2 Set the Pri DNS and Sec DNS as the Pipeline defaults.
- 3 Set 'Allow As Client DNS' to Yes or No, depending on if you want DNS information passed to users if the Client DNS information is not defined. The default for this field is Yes to permit backward compatibility. Set Allow As Client DNS to No to avoid sending the Pipeline's DNS information to users when all other IPCP DNS negotiation fails.
- 4 Select values for List Attempt and List Size.
- 5 Enter the IP address of the primary DNS server for this profile in the Client Pri DNS field.

This address is passed to a user if a DNS server is not defined in the Connection profile. It is considered not defined if set to 0.0.0.0.
- 6 Enter the IP address of the secondary DNS server for all profiles in the Client Sec DNS field.

This is the IP address of the secondary DNS server, and is the one supplied if a DNS server is not defined for the user. It is considered not defined if set to 0.0.0.0.

Configuring DNS servers in the Connection profile

To configure DNS servers in the Connection profile:

- 1 Open the IP submenu of the Connection profile.
For example:

```
30-100 Connections  
IP Options...  
LAN Adrs=0.0.0.0/0  
WAN Adrs=0.0.0.0  
IP Adrs=0.0.0.0/0  
Metric=7  
Preference=100  
Private=No  
RIP=Off  
Pool=0  
Multicast Client=No
```

```
Multicast Rate Limit=5  
Client Pri DNS=111.11.11.1  
Client Sec DNS=111.11.11.2  
Client Assign DNS=Yes
```

- 2 Enter the IP address of the primary DNS server for the dial-in user for this profile in the Client Pri DNS field.

This is the IP address that will be passed to the user when logged in using a profile. It is considered not defined if set to 0.0.0.0.

- 3 Enter the IP address of the secondary DNS server for this profile in the Client Sec DNS field.

This is the second IP address that will be passed to the user when logged in using profile. It is considered not defined if set to 0.0.0.0.

- 4 Select Yes or No for Client Assign DNS.

This value controls whether DNS information should be passed to the dial-in user or not. The default is Yes.

Local DNS host address table

You can create a local DNS table that can provide a list of IP addresses for a specific host name when the remote DNS server fails to resolve the host name successfully. The local DNS table provides the list of IP addresses only if the host name for the attempted connection matches a host name in the local DNS table.

You create the DNS table from the terminal server by entering the host names and their IP addresses in the table. A table can contain up to eight entries, with a maximum of 35 IP addresses for each entry. You enter only the first IP address; any other IP addresses in the list are automatically added if you have enabled automatic updating of the list.

You can also specify that the local DNS table is automatically updated when a connection to a host whose name matches one in the local DNS table is successfully resolved by the remote DNS. When the table is updated, the returned IP address list from the remote server replaces the stored IP addresses for that host name in the local DNS list.

You can check the list of host names and IP addresses in the table using the `termserv` command `Show Dnstab`.

Configuring the local DNS table

To enable and configure the local DNS table:

- 1 Open the Ethernet > Mod Config > DNS menu.
- 2 Select List Attempt=Yes to allow a list of the IP addresses to be displayed when using the terminal server command `Dnstab Entry`.
- 3 Select List Size and enter the number of entries you want in the list.

The minimum value is 1. The maximum value is 35.

The number of IP addresses displayed with the `Dnstab Entry` command depends upon the value you set in the List Size parameter.

If List Attempt=Yes, and the name server returns an IP address list, the list is copied into the entry in the local DNS table that matches the host name, up to the number of entries you specify in List Size. When a list of IP addresses for an entry is automatically updated, any existing list for that entry is discarded.

For example:

- If you set List Size=4 and the remote DNS returns 3 entries, the entire list of IP addresses in the local DNS table is cleared and the three returned addresses are entered for the entry.
 - If the local DNS table already contains 35 IP addresses for an entry and the remote DNS server returns only 4, or if you set List Size=4, the first four IP addresses are entered into the table for the entry and the remaining addresses in the list are set to zero.
 - If you set List Size=1, the list can contain only one IP address; any others returned by the remote DNS are ignored. If you change the List Size parameter value from a number greater than one to one, only the first IP address is retained; all others are set to zero the next time the table entry for that name is updated.
- 4 Select Enable Local DNS Table=Yes.
The default is No.
 - 5 Select Loc DNS Tab Auto Update=Yes to enable automatic updating.

The default is No. When automatic updating is enabled, the list of IP addresses for each entry is replaced with a list from the remote DNS when the remote DNS successfully resolves a connection to a host named on the table.

Creating the local DNS table

To create a local DNS table, you use the DNS table editor from the terminal server. While the editor is in use, the local DNS table is disabled for reading and updating.

Note: This procedure defines a table entry as one of the eight table indexes, which include the host name, IP address (or addresses), and information fields.

- 1 Use the DO Terminal Server command menu to open the Terminal Server. From the DO command menu, press Ctrl-D and select E-Terminal Server.

- 2 From the terminal server, enter:

```
ascend% dnstab edit
```

When the system first powers up, the table is empty. When the editor first starts up, it displays zeros for each of the eight entries in the table. To exit the table editor without making an entry, press Return.

- 3 Type an entry number and press Enter.

A warning appears if you type an invalid entry number. If the entry exists, the current name for that entry appears in the prompt.

- 4 Type the name for the current entry.

If the name is validated it is entered into the table and a prompt requests the IP address for the name that you just entered.

You can find a list of restrictions you must follow in naming entries in the DNS table at the end of this section.

- 5 Do one of the following:

Type the IP address for the entry.

The IP address is checked for format. If the format is correct, the address is entered into the table and the editor prompts for another entry.

- 6 When you are finished making entries, type O and press Return when the editor prompts you for another entry.

Editing the local DNS table

You use the DNS table editor from the terminal server to edit the DNS table entries. While the editor is in use, the local DNS table is disabled for reading and updating.

Note: This procedure defines a table entry as one of the eight table indexes, which include the host name, IP address (or addresses), and information fields.

- 1 Use the DO Terminal Server command menu to open the Terminal Server. From the DO command menu, press Ctrl-D and select E-Terminal Server.

- 2 From the terminal server, enter:

```
ascend% dnstab edit
```

If the table has already been created, the number of the entry last edited appears in the prompt.

- 3 Type an entry number or press Return to edit the entry number currently displayed.

A warning appears if you type an invalid entry number. If the entry exists, the current value for that entry appears in the prompt.

- 4 Do one of the following and press Enter.

- Type the new name for the current entry.

If the name is accepted it is entered into the table and a prompt requests the IP address for the name that you just entered.

You can find a list of restrictions you must follow in naming entries in the DNS table at the end of this section.

- Press Return to accept the current name.
- Clear the name by pressing the space bar and then Return.

If you clear an entry name and do not replace it with a new name, all information in all fields for that entry is discarded.

- 5 Do one of the following:

- If you are changing the name of the entry but not the IP address, press Return.
- To change the IP address, type the new IP address

The IP address you enter is checked for format. If the format is correct, the address is entered into the table and the editor prompts for another entry.

- 6 When you are finished making entries, type `O` and press Return when the editor prompts you for another entry.

Deleting an entry from the local DNS table

To delete an entry from the local DNS table:

- 1 Use the DO Terminal Server command menu to open the Terminal Server. From the DO command menu, press Ctrl-D and select E-Terminal Server.
- 2 To display the table, from the terminal server, enter:

```
ascend% dnstab edit
```
- 3 Type the number of the entry you want to delete and press Return.
- 4 Press the space bar and then press Return.

Restrictions for names in the local DNS table

- Names must be unique in the table.
- Names must start with an alphabetic character, either upper- or lower-case. (from A to Z or a to z).
- Names must be less than 256 characters
- Dots (periods) at the end of names are ignored.
- Names can be local names or fully qualified names that include the domain name. The Pipeline will automatically add the local domain name before it is qualified (or the secondary domain name, if the qualification with the domain name fails) from the DNS submenu of the Ethernet Profile.

User-definable TCP connection retry timeout

You can set the TCP timeout parameter to the maximum length of time the Pipeline waits to complete a connection before trying the next address supplied by a DNS server using the List Attempt feature. If the Pipeline cannot connect to the first host on the list, it tries the next, until it connects or times out.

Previously, the timeout period was not user-definable, and the timeout value was always 170 seconds, which is longer than some client software waits before

timing out. When client software timed out, the connection was dropped and no remaining addresses on the DNS list were tried. Then, each time the Pipeline restarted, it attempted the same connection that was previously unsuccessful.

To specify a timeout value, set the TCP Timeout parameter to a value from 1 to 200 seconds. Then connections to additional host addresses can be attempted before the client software times out. If the timeout value is reached and no connection is made, the Pipeline tries the next address on the list.

Setting the TCP Timeout parameter depends on the characteristics of the TCP destination hosts. For example, if the destinations are on a local network under the same administrative control as the Pipeline and are lightly loaded, then a short timeout (a few seconds) may be reasonable because a host that does not respond within that interval is probably down.

A longer timeout is appropriate if the environment includes servers with

- longer network latency times
- high loads on the net or router
- characteristics of the remote hosts are not well known

Values of 30 to 60 seconds are common in UNIX TCP implementations.

The default value, zero, specifies that the Pipeline waits for a maximum of 170 seconds to connect to each address on the list, until a connection is successful or the connection is dropped.

Network Address Translation (NAT) for a LAN

To connect to the Internet or any other TCP/IP network, a host must have an IP address that is unique within that network. The Internet and other large TCP/IP networks guarantee the uniqueness of addresses by creating central authorities that assign official IP addresses. However, many local networks use private IP addresses that are unique only on the local network. To allow a host with a private address to communicate with the Internet or another network that requires an official IP address, a Pipeline can perform a service known as network address translation (NAT). This works as follows:

- When the local host sends packets to the remote network, the Pipeline automatically translates the host's private address on the local network to an official address on the remote network.
- When the local host receives packets from the remote network, the Pipeline automatically translates the official address on the remote network to the host's private address on the local network.

NAT can be implemented to use a single address or multiple addresses. Using multiple IP addresses requires access to a remote Network Access Server (NAS) configured as a DHCP server.

Single-address NAT and port routing

A Pipeline can perform single-address NAT in these ways:

- For more than one host on the local network without borrowing IP addresses from a DHCP server on the remote network.
- When the remote network initiates the connection to the Pipeline.
- By routing packets it receives from the remote network for up to 10 different TCP or UDP ports to specific hosts and ports on the local network.

Note: You can use single-address NAT by setting the Ethernet > NAT > Lan parameter to Single IP Addr. For older units (with a switch on the back), single-address NAT is the default and the Lan parameter is hidden.

With single-address NAT, the only host on the local network that is visible to the remote network is the Pipeline.

Outgoing connection address translation

For outgoing calls, the Pipeline performs NAT for multiple hosts on the local network after getting a single IP address from the remote network during PPP negotiation.

Any number of hosts on the local network can make any number of simultaneous connections to hosts on the remote network, which is limited only to the size of the translation table. The translations between the local network and the Internet or remote network are dynamic and do not need to be preconfigured.

Incoming connection address translation

For incoming calls, the Pipeline can perform NAT for multiple hosts on the local network using its own IP address. The Pipeline routes incoming packets for up to 10 different TCP or UDP ports to specific servers on the local network.

Translations between the local network and the Internet or remote network are static and need to be preconfigured. You need to define a list of local servers and the UDP and TCP ports each would handle. You can also define a local default server that handles UDP and TCP ports not listed.

For example, you can configure the Pipeline to route all incoming packets for TCP port 80—the standard port for HTTP—to port 80 of a World Wide Web server on the local network. The port you route to does not have to be the same as the port specified in the incoming packets. For example, you can route all packets for TCP port 119, the well known port for Network News Transfer Protocol, to port 1119 on a Usenet News server on the local network. You can also specify a default server that receives any packets that aren't sent to one of the routed ports. If you don't specify any routed ports but do specify a default server, the default server receives all packets from the remote network that are sent to the Pipeline.

When you configure the Pipeline to route incoming packets for a particular TCP or UDP port to a specific server on the local network, multiple hosts on the remote network can connect to the server at the same time. The number of connections is limited by the size of the translation table.

Note: NAT automatically turns RIP off, so the address of the Pipeline is not propagated to the Internet or remote networks.

Translation table size

NAT has an internal translation table limited to 500 addresses. A translation table entry represents one TCP or UDP connection.

Note: A single application can generate many TCP and UDP connections.

The translation table entries are freed based on the following timeouts:

- Non-DNS UDP translations timeout after 5 minutes.
- DNS times out in one minute.
- TCP translations time out after 24 hours.

The translation table entries are reused as long as packets are seen that match an entry. All are freed (expired) when a connection disconnects. For Nailed connections, the connection is designed not to disconnect.

Multiple-address NAT

Multiple-address NAT can be performed when translating addresses for more than one host on the local network. To do this, the Pipeline borrows an official IP address for each host from a Dynamic Host Configuration Protocol (DHCP) server on the remote network or accessible from the remote network.

The advantages of multiple-address NAT are that hosts on the remote network can connect to specific hosts on the local network, not just specific services such as Web or FTP service, but only if the DHCP server is configured to assign the same address whenever a particular local host requests an address. Also, network service providers might require multiple-address NAT for networks with more than one host.

When you use multiple-address NAT, hosts on the remote network can connect to any of the official IP addresses that the Pipeline borrows from the DHCP server. If the local network must have more than one IP address that is visible to the remote network, you must use multiple-address NAT. If hosts on the remote network need to connect to a specific host on the local network, you can configure the DHCP server to always assign the same address when that local host requests an address.

When multiple-address NAT is enabled, the Pipeline attempts to perform IP address translation on all packets received. (It cannot distinguish between official and private addresses.)

The Pipeline acts as a DHCP client on behalf of all hosts on the LAN and relies on a DHCP server to provide addresses suitable for the remote network from its IP address pool. On the local network, the Pipeline and the hosts all have “local” addresses on the same network that are only used for local communication between the hosts and the Pipeline over the Ethernet.

When the first host on the LAN requests access to the remote network, the Pipeline gets this address through PPP negotiation. When subsequent hosts request access to the remote network, the Pipeline asks for an IP address from the DHCP server using a DHCP request packet. The server then sends an address to

the Pipeline from its IP address pool. The Pipeline uses the dynamic addresses it receives from the server to translate IP addresses on behalf of local hosts.

As packets are received on the LAN, the Pipeline determines if the source IP address has been assigned a translated address. If so, then the packet is translated, and forwarded to the wide area network. If no translation has been assigned (and is not pending), then a new DHCP request is issued for this IP address. While waiting for an IP address to be offered by the server, corresponding source packets are dropped. Similarly, for packets received from the WAN, the Pipeline checks the destination address against its table of translated addresses. If the destination address exists and is active, the Pipeline forwards the packet. If the destination address does not exist, or is not active, the packet is dropped.

IP addresses are typically offered by the DHCP server only for a limited duration, but the Pipeline automatically renews the lease on these addresses. If the connection to the remote server is dropped, all leased addresses are considered revoked. Therefore, TCP connections do not persist if the WAN call disconnects.

The Pipeline itself does not have an address on the remote network. This means that the Pipeline can only be accessed from the local network, not from the WAN. For example, you can Telnet to the Pipeline from the local network, but not from a remote network.

In some installations, the DHCP server will be handling both NAT DHCP requests and ordinary DHCP requests. In this situation, if the ordinary DHCP clients are connecting to the server over a non-bridged connection, you must have a separate DHCP server to handle the ordinary DHCP requests; the NAT DHCP server will only handle NAT DHCP requests.

Configuring single or multiple address NAT

To configure NAT on the Pipeline:

- 1 Open the menu Ethernet > NAT > NAT menu.

For example:

```
20-A01 NAT...
>Routing=Yes
Profile=NATprofile
Lan=Single IP addr
FR address=0.0.0.0
```



```
Static Mappings...
Def Server=N/A
Reuse last addr=N/A
Reuse addr timeout=N/A
```

- 2 Enable NAT by setting Routing to Yes. Without this setting, no other setting is valid.
- 3 Set Profile to the name of a Connection profile you want to use to connect to the Network Access Server (NAS).
- 4 The Lan parameter can be set to Single IP Addr (by default) or to Multiple.
- 5 FR address refers to Frame Relay. Refer to “NAT for Frame Relay” on page 3-28 for more information.
- 6 The Static Mappings menu includes 10 Static Mapping *nn* submenus, where *nn* is a value from 01 to 10. Each of these submenus contains parameters for controlling the translation of the private IP addresses to TCP or UDP port numbers when operating in single-address NAT mode. You only need to specify static mappings for connections initiated by devices calling into the private LAN. For sessions initiated by hosts on the private LAN, the Pipeline generates a mapping dynamically if one does not already exist in the Static Mappings parameters.

Each Static Mapping *nn* menu contains the following parameters:

```
20-A00 NAT
  Static Mapping 01
    Valid=Yes
    Dst Port#=21
    Protocol=TCP
    Loc Port#=21
    Loc Adrs=181.100.100.102
```

See “Routing incoming sessions for up to 10 servers on a LAN” on page 3-30 for information about how to set each parameter.

- 7 Optionally set Def Server to the IP address of a local server to which the Pipeline routes incoming packets that are *not* routed to a specific server and port. (See “Routing all incoming sessions to the default server” on page 3-29 for more information.)

- 8 Optionally set Reuse last addr to Yes to continue to use a dynamically assigned IP address. The Reuse addr timeout value specifies the time to use the address. Set it to a number of minutes (up to 1440). Limitations apply, which are described in the *Reference Guide*.
- 9 Exit and save the profile.

Note: If you have additional routers on your local area network, open Ethernet > Mod Config > Ether Options, and set the value of Ignore Def Rt to Yes. This avoids the possibility that a default route from the ISP will overwrite the NAT route.

NAT for Frame Relay

The single-IP address implementation of NAT extends to Frame Relay. Connections using Frame Relay encapsulation to the Pipeline running single-IP address NAT, translate the local addresses into a single, official address set by the FR address parameter.

Set the Routing parameter in the NAT profile to enable NAT. Set the Lan parameter to Single IP addr.

```
20-A00 NAT
20-A01 NAT...
Routing=Yes
Profile=max4
Lan=Single IP addr
FR address=0.0.0.0
Static Mapping...
Def Server=181.81.8.1
Reuse last addr=No
Reuse addr timeout=N/A
```

When Routing=Yes and a valid, official IP address is entered for FR address, NAT is enabled for Frame Relay connections.

Configuring NAT port routing (Static Mapping submenu)

The Static Mappings menu includes 10 Static Mapping *nn* submenus, where *nn* is a value from 01 to 10. Each of these submenus contains parameters for controlling the translation of a private IP address and port number to a TCP or UDP port number. Static Mappings applies only to single-address NAT. You only need to specify static mappings for connections initiated by devices calling into the private LAN.

You can configure a NAT port routing

- to define a default server on the local private LAN
The Pipeline routes incoming packets to the default server when their destination port number does not match an entry in Static Mappings nor does it match a port number dynamically assigned when a local host initiates a TCP / UDP session.
- to define a list of up to 10 servers & services on the local private LAN
The Pipeline routes incoming packets to hosts on the local private LAN when their destination port matches one of the 10 destination ports in Static Mappings.

Note: You need to configure port routing only for sessions initiated by hosts outside the private LAN. For sessions initiated by hosts on the private LAN, the Pipeline generates the port mapping dynamically.

For port routing in single-address NAT to work, if firewalls are present, they must be configured to allow the Pipeline to receive packets for the routed ports.

Routing all incoming sessions to the default server

To configure the Pipeline to perform NAT and to define a single server which handles all sessions initiated by callers from outside the private LAN:

- 1 Open the Ethernet > NAT > NAT menu.
- 2 Set the Routing parameter to Yes.
- 3 Set the Profile parameter to the name of an existing Connection profile.

The Pipeline performs NAT whenever a connection is made with this Connection profile. The connection can be initiated either by the Pipeline or by the remote network.

- 4 Set the Lan parameter to Single IP Addr.
- 5 If you previously configured the Pipeline to route incoming packets for specific TCP or UDP ports (as described in “Routing incoming sessions for up to 10 servers on a LAN” on page 3-30).
 - Open each Ethernet > NAT > Static Mapping > Static Mapping *nn* menu (where *nn* is a number between 01 and 10).
 - Set the Valid parameter in each menu to No.
- 6 Set the Def Server parameter to the IP address of the server on the local network to receive all incoming packets from the remote network.
- 7 Press the Esc key to exit the menu.
- 8 Save the changes when prompted.

The changes take effect the next time a connection is made for the NAT profile. To make the changes immediately, close the connection specified by the Profile parameter and then reopen it.

Routing incoming sessions for up to 10 servers on a LAN

To configure the Pipeline to perform NAT and to define up to 10 servers and optionally a default server which handle sessions initiated by callers from outside the private LAN:

- 1 Open the Ethernet > NAT > NAT menu.
- 2 Set the Routing parameter to Yes.
- 3 Set the Profile parameter to the name of an existing Connection profile.

The Pipeline performs NAT whenever a connection is made with this Connection profile. The connection can be initiated either by the Pipeline or by the remote network.
- 4 Set the Lan parameter to Single IP Addr.
- 5 Open the Ethernet > NAT > NAT > Static Mapping menu.
- 6 Open a Static Mapping *nn* menu, where *nn* is a number between 01 and 10.

You use the parameters in each Static Mapping *nn* menu to specify routing for incoming packets sent to a particular TCP or UDP port.

- 7** Set the Valid parameter to Yes.

This enables the port routing specified by the remaining parameters in the menu. Setting this parameter to No disables routing for the specified port.

- 8** Set the Dst Port # parameter to the number of a TCP or UDP port which users outside the private network can access. Each Dst Port # corresponds to a service provided by a server on the local private network. You can use the actual port number as given by the Loc Port # parameter as long as that address is unique for the local private network. See “Well-known ports” on page 3-32 for information on obtaining port numbers.

The Pipeline routes incoming packets it receives from the remote network for this port to the local server and port you’re about to specify.

- 9** Set the Protocol parameter to TCP or UDP.

This parameter determines whether the Dst Port # and Loc Port # parameters specify TCP ports or UDP ports.

- 10** Set the Loc Port # to a port corresponding to a service provided by the local servers.

- 11** Set the Loc Adrs parameter to the address of the local server providing the service specified by Loc Port #.

- 12** Exit and save the profile.

Repeat steps 6 through 12 for any additional ports whose packets you want to route to a specific server and port on the local network.

- 13** Open the Ethernet > NAT > NAT menu.

- 14** Set the Def Server parameter to the IP address of a server on the local network that receives any remaining incoming packets from the remote network, that is, any that aren’t for ports you’ve specified in Static Mapping *nn* menus.

- 15** Exit and save the profile.

The changes take effect the next time a connection is made for the NAT Profile. To make the changes immediately, close the connection specified by the Profile parameter and then reopen it.

Disabling routing for specific ports

To disable routing of incoming packets from a remote network for specific TCP or UDP ports:

- 1** Open the Ethernet > NAT > NAT > Static Mapping menu.
- 2** Open a Static Mapping *nn* menu, where *nn* is a number between 01 and 10. The parameters in each Static Mapping *nn* menu specify the routing for incoming packets sent to a particular TCP or UDP port.
- 3** Set the Valid parameter to No. This disables routing for the port specified by the Dst Port# and Protocol parameters in this menu.
- 4** Exit and save the profile. Repeat steps 2 through 4 to disable routing for any additional ports.
- 5** Exit and save the profile.

The changes take effect the next time a connection is made for the NAT Profile. To make the changes immediately, close the connection specified by the Profile parameter and then reopen it.

Well-known ports

TCP and UDP ports numbered 0-1023 are called Well Known Ports. These ports, which include the ports for the most common services available on the Internet, are assigned by the Internet Assigned Numbers Authority (IANA). In almost all cases, the TCP and UDP port numbers for a service are the same.

You can obtain current lists of Well Known Ports and Registered Ports (ports in the range 1024-4915 that have been registered with the IANA) via FTP from `ftp://ftp.isi.edu/in-notes/iana/assignments/port-numbers`

Configuring IPX Routing

This chapter includes the following topics:

How the Pipeline performs IPX routing	4-1
Adding the Pipeline to the local IPX network	4-12
Working with the RIP and SAP tables	4-15
Configuring IPX routing connections	4-24

How the Pipeline performs IPX routing

To support Internet Packet Exchange (IPX) routing between sites that run Novell NetWare version 3.11 or later, the Pipeline operates as an IPX router, with one interface on the local Ethernet and the other across the wide-area network (WAN). Each IPX Connection profile is an IPX WAN interface.

The most common uses for IPX routing in the Pipeline are to:

- Integrate multiple NetWare local-area networks (LANs) to form an interconnected WAN.
- Allow dial-in NetWare clients to access local NetWare services.

The Pipeline supports IPX routing over Point-to-Point Protocol (PPP), Multilink PPP (MP), and frame relay connections. Support for both the IPXWAN and PPP Internet Protocol Control Protocol for IPX (IPXCP) makes the Pipeline fully

interoperable with other vendors' products that conform to these protocols and associated RFCs.

Note: IPX can be transmitted using different frame types. The Pipeline routes only one IPX frame type, and it routes and spoofs IPX packets only if they are encapsulated in that type of frame. If bridging is enabled in the same Connection profile as IPX routing, the Pipeline will bridge any other IPX-packet frame types. (For more information see Chapter 5, "Configuring the Pipeline as a Bridge.")

Unlike an IP routing configuration, where the Pipeline uniquely identifies the calling device by its IP address, an IPX routing configuration does not include a built-in way to uniquely identify callers. For that reason, password authentication is required unless IP routing is configured in the same Connection profile. (For details, see Chapter 7, "Setting Up Pipeline Security.")

IPX Service Advertising Protocol (SAP) tables

The Pipeline follows standard IPX SAP behavior for routers. However, when the connection is to another Pipeline configured for IPX routing, both ends of the connection exchange their entire SAP tables, so all remote services are immediately added to each unit's SAP table.

NetWare servers broadcast SAP packets every 60 seconds to make sure that routers know about their services. Routers build a SAP table with an entry for each service advertised by each known server. When a router stops receiving SAP broadcasts from a server, it ages the SAP-table entry and eventually removes it from the table.

Routers use SAP tables to respond to client queries. When a NetWare client sends a SAP request to locate a service, the Pipeline consults its SAP table and replies with its own hardware address and the internal address of the requested server (similar to "Enabling proxy mode in the Pipeline" on page 3-6).

The client can then transmit packets whose destination address is the internal address of the server. When the Pipeline receives those packets, it consults its RIP table. If it finds an entry for that destination address, it brings up the connection or forwards the packet across the active connection.

IPX Routing Information Protocol (RIP) tables

IPX RIP is similar to the routing information protocol in the TCP/IP protocol suite, but it is a different protocol. In this chapter, RIP always refers to IPX RIP.

The Pipeline follows standard IPX RIP behavior for routers when connecting to other-vendor units. However, when it connects to another Pipeline configured for IPX routing, both ends of the connection immediately exchange their entire RIP tables. In addition, the Pipeline maintains those RIP entries as static until the unit is reset or power-cycled.

The destination of an IPX route is the internal network of a server. For example, NetWare file servers are assigned an internal IPX network number by the network administrator and typically use the default node address of 000000000001. This is the destination network address for file read/write requests. (If you are not familiar with internal network numbers, see your NetWare documentation for more information.)

IPX routers broadcast RIP updates periodically and whenever a WAN connection is established. The Pipeline receives RIP broadcasts from a remote device, adds 1 to the hop count of each advertised route, updates its own RIP table, and broadcasts updated RIP packets on connected networks in a split-horizon fashion.

The Pipeline recognizes network number -2 (FFFFFFFE hex) as the IPX RIP default route, and forwards any packet with an unrecognized address to the IPX router advertising that default route. For example, if the Pipeline receives an IPX packet destined for network 77777777 and it does not have a RIP table entry for that destination, the Pipeline forwards the packet towards network number FFFFFFFE, if available, instead of simply dropping the packet. If more than one IPX router is advertising the default route, the Pipeline bases its routing decision on Hop and Tick count.

Extensions to standard IPX

NetWare uses dynamic routing and service location to let clients locate a server dynamically, regardless of where it is physically located. This scheme is designed for LAN environments. For WAN functionality, the Pipeline provides the following extensions to standard IPX:

- Virtual IPX network defined for dial-in clients

Configuring IPX Routing

How the Pipeline performs IPX routing

- IPX Route profiles
- IPX SAP filters
- Dial Query
- Watchdog spoofing

Virtual IPX network for dial-in clients

The Pipeline allows individual NetWare clients that do not have an IPX network address to use an IPX routing connection to the local network if they are running PPP client software.

To enable the Pipeline to route to such dial-in clients, you must specify an IPX network number in the Ethernet profile. The number must be unique within the entire IPX routing domain of the Pipeline (the local routing domain as well as all WAN links). It defines a “virtual” IPX network reserved for dial-in clients.

The Connection profile for each dial-in client must specify “Dialin” for the Pipeline to assign the virtual IPX network number to the dial-in client during PPP negotiation. If the client does not provide its own unique node number, the Pipeline assigns a node number as well as the network number. It does not send RIP and SAP advertisements across the connection, and ignores RIP and SAP advertisements received from the far end. However, it does respond to RIP and SAP queries received from dial-in clients.

For more information, see “Defining a virtual IPX network for dial-in clients” on page 4-15 and “An example dial-in client connection” on page 4-24.

Optimized access for dial-in NetWare clients

Without optimized access, the Pipeline assumes that the far end of an incoming IPX connection is another IPX router. After answering the call, the Pipeline recognizes the caller as a client via the Peer=Dialin setting in the caller’s Connection profile.

The Answer profile also contains a Peer parameter to enable the Pipeline to treat incoming IPX connections as clients even when configured profiles are not in use. You must set this for dial-in Windows 95 clients with no configured profile, because without it, the connection can take more than a minute to establish and the client cannot see NetWare servers on the local network.

The IPX Options submenu in the Answer profile contains the Peer parameter which enables the Pipeline to route to dial-in NetWare clients when the client has no configured profile. The Peer parameter is set to Router by default, which tells the Pipeline to negotiate inbound IPX calls as if the far end is a router. The Dialin setting tells the Pipeline to negotiate inbound IPX calls as if the far end is a dial-in NetWare client.

The following list shows the Peer parameter as well as other required parameters:

```
Answer
  Profile Req=No
  IPX options...
    Peer=Dialin
  PPP options...
    Route IPX=Yes

Mod Config
  Ether options...
    IPX Enet#=cffff123
    IPX Pool#=cf000888
```

Required settings

When enabling this feature, consider the following:

- Calls for which no Connection profile is found must be answered.
The call might require authentication, or use SecureID passwords.
The dial-in client must be running PPP software.
- IPX routing must be enabled in the PPP Options submenu of the Answer profile, and the IPX network number of the router's Ethernet interface must be configured in the Ethernet profile.
- Specify an IPX Pool number in the Ethernet profile, so that the Pipeline can route to dial-in clients.

The network number must be unique within the entire IPX routing domain of the Pipeline (the local routing domain as well as all WAN links). This is a “virtual” IPX network reserved for dial-in clients. If the client does not provide its own unique node number, the Pipeline assigns a unique node number to the client as well.

Note: The Pipeline does not send RIP and SAP advertisements across the connection and ignores RIP and SAP advertisements received from the far

Configuring IPX Routing

How the Pipeline performs IPX routing

end. However, it does respond to RIP and SAP queries received from dial-in clients.

IPX Route profiles

Static IPX routes are specified in IPX Route profiles. When the Pipeline unit's RIP and SAP tables are cleared due to a reset or power-cycle, the static routes are added when the unit initializes. Each static route contains the information needed to reach one server.

When the Pipeline is connecting to another Pipeline, you can choose not to configure a static route. Instead, you can use the DO menu to manually dial the initial connection to that site following a power-cycle or reset. Once connected, the Pipeline downloads the RIP table from the remote site and maintains the routes as static until the next power-cycle or reset.

Static routes need manual updating whenever the specified server is removed or has an address change. However, static routes are a way to ensure that the Pipeline can bring up the appropriate connection in response to clients' SAP requests and to prevent timeouts when a client takes a long time to locate a server on the WAN. (For more information, see "Configuring a static IPX route" on page 4-18.)

You can also specify a route to a destination IPX network without defining an IPX server in the IPX Routes profile. You can reach an IPX network by entering the Network number (for example, Network=00123456) without specifying the Server Name and Server Type.

To configure IPX routes, open Ethernet > IPX Routes > *any profile*

For example:

```
Server Name=server-name
Active=Yes
Network=CC1234FF
Node=000000000001
Socket=0000
Server Type=0004
Hop Count=2
```

```
Tick Count=12
Connection #=0
```

Note: The Pipeline cannot support more than 300 server and route entries. In order to keep the Pipeline operational with IPX enabled on a large network, the Pipeline enforces a maximum limit of 300 server and route entries, including limit checking for both server and route entries. When the Pipeline reaches its limit of 300, it drops all IPX route and SAP packets containing additional routes and services. This limit results in an incomplete network map, so you need to activate a size-limiting feature, such as enabling IPX SAP Proxy or IPX filtering. (To check the number of current servers and routes, see “Using the terminal server interface” on page 8-17, specifically note the function of the `show netw servers` command. For information about how to use the IPX SAP Proxy parameter, refer to the *Reference Guide*, and also see “IPX SAP proxy servers” on page 4-14. For information on setting up IPX filtering, refer to “Managing IPX SAP filters” on page 4-21.)

IPX SAP filters

You might not want the Pipeline SAP table to include long lists of all servers available at a remote site. IPX SAP filters let you exclude services from the SAP table, or explicitly include certain services.

SAP filters can be applied to inbound or outbound SAP packets. Inbound filters control which services are added to the Pipeline unit’s SAP table from advertisements on a network link. Outbound filters control which services the Pipeline advertises on a particular network link. (For more information, see “Managing IPX SAP filters” on page 4-21.)

Configure IPX SAP filters in Ethernet > IPX SAP filters > *any profile*:

```
Name=optional
Input SAP filters...
Output SAP filters
    Valid=Yes
    Type=Exclude
    Server Type=0004
    Server Name=SERVER-1
```

See “Managing IPX SAP filters” on page 4-21 for details on each parameter.

IPX Type 20 packet propagation support

Some applications, such as NetBIOS, use IPX Type 20 packets to broadcast names over a network. By default, these broadcasts are not propagated over routed links (as Novell recommends), and are not forwarded over links that have less than 1 Mbps throughput.

Since the Pipeline cannot support these types of applications, you can change the setting of IPX Type 20 packet propagation to Yes if required.

To support IPX Type 20 propagation:

- 1 Open Ethernet > Mod Config > Ether Options
- 2 Set Handle IPX Type20 to Yes.

Dial Query

Dial Query is a Connection profile parameter that instructs the Pipeline to bring up that connection when it receives a SAP query for service type 0004 (a file server) when that service type is not present in the Pipeline SAP table. If the Pipeline has no SAP table entry for service type 0004, it brings up every connection that has Dial Query set to Yes. For example, if five Connection profiles have Dial Query set to Yes, the Pipeline brings up all five connections in response to the query.

Note: If the Pipeline has a static IPX route to a remote server, it will bring up that connection instead of the more costly solution of bringing up every connection that has Dial Query set.

To configure Dial Query, open Ethernet > Connection > *any profile* > IPX Options:

```
Peer=Dialin (used for dynamic addressing)
IPX RIP=None
IPX SAP=Send
Dial Query=No
Handle IPX=Client (used for IPX client bridging)
Netware t/o=30 (watchdog spoofing)
```

For information about the Handle IPX parameter and IPX bridging, see Chapter 5, “Configuring the Pipeline as a Bridge.”

Watchdog spoofing

NetWare servers send out NCP watchdog packets to monitor client connections. Clients that respond to watchdog packets remain logged into the server. If a client does not respond to watchdog packets for a certain amount of time, the server logs the client out.

Repeated watchdog packets can cause a WAN connection to stay active. But if the Pipeline filters out the packets, client logins are dropped by the remote server. To prevent repeated client logouts while allowing WAN connections to be brought down in times of inactivity, the Pipeline responds to watchdog requests as a proxy for remote IPX routed or bridged clients. Responding to NCP requests is commonly called watchdog spoofing. To the server, a spoofed connection looks like a normal, active client login session, so it does not log the client out.

When a remote client link goes down, the timer begins counting. When the value of the Netware t/o (timeout) field is reached, the Pipeline stops responding to watchdog packets for the client, and the connection is released by the server. If there is a reconnection of the WAN session before the timeout value is reached, the timer is reset.

Note: The Pipeline software filters IPX watchdog packets automatically on all IPX routing connections and all IPX bridging connections that have watchdog spoofing enabled. The Pipeline applies a call filter implicitly, which prevents the idle timer from resetting when IPX watchdog packets are sent or received. This filter is applied after the standard data and call filters.

Automatic SPX spoofing

NetWare applications that require a guaranteed packet delivery use the NetWare SPX protocol. This includes applications such as Print Server (PSERVER) and Remote Printer (RPRINTER), as well as Remote Console (RCONSOLE). The client's SPX watchdog monitors the connection with the server while the connection is idle. To monitor the connection, the SPX watchdog sends a query that brings up the WAN connection every 14 seconds while an SPX application is running.

The Pipeline lets Netware SPX clients stay logged in without keeping the WAN connection up in times of inactivity by automatically responding to SPX watchdog requests from the LAN with a spoofed SPX-watchdog-reply packet,

and drops any SPX-watchdog keep-alive packets from the LAN, without sending them on to the WAN. You do not need to set any parameters to enable this function; however, note that routers on both ends of the connection must support this feature for it to function.

WAN considerations for NetWare client software

In most cases, NetWare clients on a wide-area network do not need special configuration. But the following issues sometimes affect NetWare clients in an IPX routing environment:

- Preferred servers

If the local IPX network supports NetWare servers, configure NetWare clients with a preferred server on the local network, not at a remote site. If the local Ethernet does not support NetWare servers, configure local clients with a preferred server on the network that requires the least expensive connection costs. (For more information, see your NetWare documentation.)

- Local copy of LOGIN.EXE

Due to possible performance issues, executing programs remotely is not recommended. You should put LOGIN.EXE on each client's local drive.

- Packet Burst (NetWare 3.11)

Packet Burst lets servers send a data stream across the WAN before a client sends an acknowledgment. It is included automatically in server and client software for NetWare 3.12 or later. If local servers are running NetWare 3.11, they should have PBURST.NLM loaded. (For more information, see your NetWare documentation.)

- Macintosh or UNIX clients

Both Macintosh and UNIX clients can use IPX to communicate with servers. However, both types of clients have native support for AppleTalk (Macintosh) or TCP/IP (UNIX).

If Macintosh clients need to access NetWare servers across the WAN using AppleTalk (rather than MacIPX), the WAN link must support bridging, or else the AppleTalk packets will not make it across the connection.

If UNIX clients need to access NetWare servers using TCP/IP (rather than UNIXWare), the Pipeline must be configured as a bridge or IP router, or else the TCP/IP packets will not make it across the connection.

IPX in the Answer profile

Before the Pipeline answers an incoming call, it checks the settings in its Answer profile. If the call does not include the information required by the Answer profile, the Pipeline hangs up.

Note: Unlike an IP routing configuration, where the Pipeline uniquely identifies the calling device by its IP address, an IPX routing configuration does not include a built-in way to uniquely identify callers. For that reason, password authentication is required unless IP routing is configured in the same Connection profile.

To set the Answer profile parameters that enable incoming IPX routing calls:

- 1 Open the Ethernet > Answer > PPP Options menu.
- 2 Turn on IPX routing:
`Route IPX=Yes`
- 3 Turn on authentication.
For example:
`Recv Auth=Either`
For more information about setting up password authentication, see Chapter 7, “Setting Up Pipeline Security.”

To apply an IPX SAP filter profile to the Answer profile:

- 4 Open the Ethernet > Answer > Session options submenu.
- 5 Specify the number of the IPX SAP filter profile you have defined.
You apply an IPX SAP filter profile by specifying the unique part of the number it is assigned in the IPX SAP Filters menu (such as 1, 2, 3,...). For example:
`IPX SAP Filter=1`
For details, see “Managing IPX SAP filters” on page 4-21.
- 6 Close the Answer profile.

Adding the Pipeline to the local IPX network

To connect the Pipeline to your local IPX network, you must perform the following tasks:

- Turn on IPX routing.
- Specify the IPX frame type the Pipeline will route and watchdog spoof.
- Specify the Pipeline IPX network number (or allow it to learn the number from other routers).

In addition, you might want to define an IPX network number for dial-in clients.

Checking local NetWare configurations

IPX packets are supported in more than one Ethernet frame type on an Ethernet segment. However, the Pipeline can only route and perform watchdog spoofing for the IPX frame type you specify. (It will bridge other IPX packet types if bridging is enabled.)

To check the IPX configuration of a NetWare server on the local Ethernet:

- 1 Go to the NetWare server's console.
- 2 Type `LOAD INSTALL` to view the `AUTOEXEC.NCF` file.
- 3 Look for lines similar to these:

```
internal network 1234
Bind ipx ipx-card net=CF0123FF
Load 3c509 name=ipx-card frame=ETHERNET_8023
```

The first line specifies the internal network number of the server. If you are not familiar with internal network numbers, see your NetWare documentation. The Pipeline does not require internal network numbers.

The “Bind” line specifies the IPX network number in use on the Ethernet. The Pipeline must use the same IPX network number for its Ethernet interface. You can specify the number explicitly in the Pipeline Ethernet profile, or leave the Pipeline number set to zero to enable it to “learn” the number from other routers.

The “Load” line specifies the packet frame being used by this server’s Ethernet controller (in this example, 802.3 frames). If you are not familiar with the concept of packet frames, see your NetWare documentation.

Note: IPX network numbers on each network segment, and internal network within any server, on the *entire WAN* must each have a unique network number. So you should know the external and internal network numbers in use at all sites.

Configuring IPX on the Pipeline Ethernet interface

By default, when you turn on IPX routing in the Pipeline and close the Ethernet profile, the Pipeline comes up in IPX routing mode, uses the default frame type 802.2 (which is the suggested frame type for NetWare 3.12 or later), and listens on the Ethernet to acquire its IPX network number from other IPX routers on that segment.

To turn on IPX routing in the Pipeline:

- 1 Open the Ethernet > Mod Config profile.
- 2 Turn on IPX routing:
`IPX Routing=Yes`

To specify the IPX frame type:

- 1 Open Ethernet > Mod Config > Ether Options.
- 2 Select the IPX frame type.
For example:
`IPX Frame=802.2`

Note: Make sure that the type you choose is consistent with the frame type in use by most servers on the local network.

To allow the Pipeline to learn its IPX network number:

- 1 Set the IPX Enet number to zero.
`IPX Enet #=00000000`

This causes the Pipeline to listen for its network number and acquire it from another router. Or you can enter an IPX network number other than zero, for example:

Configuring IPX Routing

Adding the Pipeline to the local IPX network

```
IPX Enet #=C90AB997
```

Note: If you specify an IPX network number other than zero, the Pipeline becomes a “seeding” router and other routers can learn their number from the Pipeline. In that case, make sure that the number you enter is the same one used by other IPX routers on the same network. (For more information about seeding routers, see your NetWare documentation.)

- 2 Close and save the Ethernet > Mod Config profile.

You can IPXPing the Pipeline from a NetWare server or client to verify that it has acquired its IPX address and is up and running on the network.

IPX SAP proxy servers

Some networks are designed to prevent the propagation of RIP and SAP packets. The IPX SAP proxy parameter lets you point to an IPX SAP proxy server. To ensure that remote users can connect, there are three default IPX SAP proxy servers in the Ethernet > Mod Config > Ether options menu.

Using IPXping to check the configuration

The IPXping command enables you to verify the transmission path to NetWare stations at the network layer. It works on the same LAN as the Pipeline or across a WAN connection that has IPX routing enabled.

Enter the IPXping command in this format:

```
ipxping hostname
```

where hostname is either the IPX address of the NetWare workstation or the advertised name of a server. The IPX address consists of the IPX network and node numbers for a station, as in:

```
ipxping CFFF1234:000000000001
```

If you are using IPXping to verify connectivity with an advertised NetWare server, you can simply enter the name of the server, as in:

```
ipxping server-1
```

You can terminate the IPXping at any time by pressing Ctrl-C.

Defining a virtual IPX network for dial-in clients

Dial-in clients do not belong to an IPX network, so they must be assigned an IPX network number to establish a routing connection with the Pipeline. To provide an IPX network number for dial-in clients, you must define a virtual IPX network in the Ethernet profile. The Pipeline advertises the route to this virtual network and assigns it as the network address for dial-in clients.

Note: The most common configuration mistake on NetWare internetworks is in assigning duplicate network numbers. Make sure that the network number you specify in the IPX Pool# field is unique within the entire IPX routing domain of the Pipeline unit.

To configure the Pipeline with an IPX network for dial-in clients:

- 1 Open the Ethernet > Mod Config > Ether options menu.
- 2 Set the IPX Pool # parameter to a 32-bit hexadecimal IPX network number that is unique within your entire IPX routing domain.

For example:

```
IPX Pool #=cccc1234
```

- 3 Close the Ethernet profile.

Working with the RIP and SAP tables

In managing the RIP and SAP tables, you might want to perform one or more of the following tasks:

- View the RIP and SAP tables.
- Configure RIP in a Connection profile.
- Configure a static route.
- Configure SAP in a Connection profile.
- Define and apply an IPX SAP filter.

Discussion about performing each of these tasks follows. Additionally, you might want to define standard call filters or data filters to control WAN traffic and connections. Call and data filters are discussed in Chapter 6, “Defining Filters and Firewalls.”

Viewing the RIP and SAP tables

To see the current RIP table, invoke the terminal server (described on page 8-17) and type:

```
show netware networks
```

The current RIP table will be displayed, and will be similar to the following:

network	next router	hops	ticks	origin
22222222	0000000000000	2	12	nov12-m2 S
A30E0A04	0080A30E0A04	1	3	Ethernet
A30E1347	0080A30E1347	1	3	Ethernet
A30E0EB8	0080A30E0EB8	1	3	Ethernet
A304B294	0080A304B294	1	3	Ethernet
EE000001	00608CB24081	1	3	Ethernet
AA000002	0000000000000	0	1	Ethernet S

The RIP table includes these fields:

- Network. Internal network number of a NetWare server.
- Next Router. Address of an IPX router used to forward packets to that server.
- Hops. Hop count to the destination network (server).
- Ticks. Tick count (18 ticks/second) to the destination network (server).
Best routes are calculated on the basis of tick count, not hop count.
- Origin. Name of the Connection profile used to reach the server.

To see the current IPX SAP table, in the terminal server, type the following:

```
show netware servers
```

You'll see a SAP table similar to the following:

IPX address	type	server name
EE000001:000000000001:0040	026b	SERVER1__
EE000001:000000000001:4510	0004	NOVL1
EE000001:000000000001:4005	0278	SERVER2__
A30E0A04:000000000001:8060	0047	EPS_0E0A04
A30E1347:000000000001:8060	0047	EPS_0E1347
A30E0EB8:000000000001:8060	0047	EPS_0E0EB8
A30EB294:000000000001:8060	0047	EPS_04B294

Fields in the SAP table, and their contents, are:

- **IPX Address.** IPX address of one server.
The IPX address uses the following format:
network number:node number:socket number
- **Service Type.** Hexadecimal value representing a type of NetWare service.
For example, the number for file servers is 0004.
- **Server Name.** Server's name (up to 35 characters).

Configuring RIP in a Connection profile

By default, the IPX RIP parameter in a Connection profile is set to Both, indicating that RIP broadcasts will be exchanged in both directions. You can disable the exchange of RIP broadcasts across a WAN connection, or specify that the Pipeline will only send or only receive RIP broadcasts on that connection. (If the Pipeline does not receive RIP broadcasts from a remote unit, you should configure a static route to at least one server on that network. See "Configuring a static IPX route" on page 4-18.)

To restrict RIP exchanges across a WAN connection:

- 1 Open a Connection profile that has IPX routing enabled.
- 2 Open the IPX Options submenu.
- 3 Set the IPX RIP parameter to a value other than the default setting of Both.
For example:

IPX RIP=Recv

This setting specifies that the Pipeline receives the RIP table from the other IPX router but will not upload its RIP table. To disable IPX RIP, set:

IPX RIP=None

- 4 Close the Connection profile.

Configuring a static IPX route

Each static IPX route contains all of the information needed to reach one NetWare server on a remote network. When the Pipeline receives an outbound packet for that server, it finds the referenced Connection profile and dials the connection.

Note: You don't need to create IPX routes to servers on the local Ethernet.

Most sites configure only a few IPX routes and rely on RIP for most other connections. If you have servers on both sides of the WAN connection, you should define a static route to the remote site even if your environment requires dynamic routes. If you have one static route to a remote site, it should specify a "master" NetWare server that knows about many other services. NetWare workstations can then learn about other remote services by connecting to that remote NetWare server.

Note: Remember that static IPX routes are manually administered, so they must be updated if there is a change to a remote server.

To define an IPX Route profile:

- 1 Open Ethernet > IPX Routes > *any profile*.

For example:

```
Server Name=SERVER-1
Active=Yes
Network=ccccfff1
Node=000000000001
Socket=0000
Server Type=0004
Hop Count=2
Tick Count=12
Connection #=1
```

- 2 Specify the name of the remote NetWare server.

For example:

Server Name=SERVER-1

- 3** Specify that the route should be added to the RIP table:

Active-Yes

- 4** Enter the remote server's internal network number.

For example:

Network=ABC01FFF

- 5** Enter the remote server's node number.

For example:

Node=0000000000001

The default 0000000000001 is typically the node number for NetWare file servers.

- 6** Specify the remote server's socket number.

For example:

Socket=0451

Typically, Novell file servers use socket 0451.

The number you specify must be a well-known socket number. Services that use dynamic socket numbers might use a different socket each time they load and will not work in IPX Route profiles. To bring up a connection to a remote service that uses a dynamic socket number, specify a "master" server with a well-known socket number on that network.

- 7** Specify the SAP Service Type.

For example:

Service Type=0004

NetWare file servers are SAP Service type 0004.

- 8** Specify the distance in hops to the server.

For example:

Hop count=2

Usually the default of 2 is appropriate.

- 9** Specify the distance to the server in ticks (18 ticks/second).

For example:

Tick count=12

Usually the default of 12 is appropriate, but you might need to increase this value for very distant servers.

- 10** Specify the number of the Connection profile that defines the WAN connection.

A Connection profile is referenced by the unique part of the number it is assigned in the Connections menu (1, 2, 3, and so forth).

```
Connection #=2
```

- 11** Close the IPX Route profile.

Configuring SAP in a Connection profile

By default, the IPX SAP parameter in a Connection profile is set to Both, indicating that SAP broadcasts will be exchanged in both directions. If SAP is enabled to both send and receive broadcasts on the WAN interface, the Pipeline broadcasts its SAP table to the remote network and listens for service updates from that network. Eventually, both networks have a table of all services on the WAN.

To control which services are advertised and where, you can disable the exchange of SAP broadcasts across a WAN connection, or specify that the Pipeline will only send or only receive SAP broadcasts on that connection.

To restrict SAP broadcasts across a WAN connection:

- 1** Open a Connection profile that has IPX routing enabled.
- 2** Open the IPX Options submenu.
- 3** Set the IPX RIP parameter to a value other than the default setting of Both.

For example:

```
IPX SAP=Recv
```

This setting specifies that the Pipeline receives SAP table updates from the remote router. If you do not want the Pipeline to send or receive SAP broadcasts on this connection, use the following setting:

```
IPX SAP=None
```

- 4** Close the Connection profile.

Managing IPX SAP filters

IPX SAP filters include or exclude specific NetWare services from the Pipeline unit's SAP table.

Note: IPX SAP filters control which services are added to the local SAP table or passed on in SAP response packets across IPX routing connections (*not* IPX bridging connections). IPX SAP filters are used to manage connectivity costs, unlike filters that prevent periodic RIP and SAP broadcasts from keeping a connection up unnecessarily.

Defining an IPX SAP filter

To define an IPX SAP filter:

- 1 Open Ethernet > IPX SAP Filter > *any profile*.

For example:

```
Name=optional
Input filters...
Output filters...
    Valid=Yes
    Type=Exclude
    Server Type=0004
    Server Name=SERVER-5
```

- 2 Specify a name for the profile.

- 3 Open the list of Input filters.

Input filter conditions are applied to all SAP packets received by the Pipeline. They screen advertised services.

You can specify up to 12 filters to include or exclude services from particular servers. These filters are applied in the order listed in the Input Filters menu.

```
Filter name
    In filter 01
    Valid=Yes
    Type=IPX
    Generic...
    Ip...
    Ipx...
```

Configuring IPX Routing

Working with the RIP and SAP tables

When the IPX filter type is specified, the following IPX submenu is available:

```
Ipx...  
  Forward=No  
  Src Network Adrs=cfff0000  
  Dst Network Adrs=cf088888  
  Src Node Adrs=111222333  
  Dst Node Adrs=aaabbbccc  
  Src Socket Cmp=equal  
  Src Socket #=0451  
  Dst Socket Cmp=equal  
  Dst Socket #=0015
```

The Forward parameter works just as it does for other filter types. If it is set to No, a matching packet is discarded. The following new filter parameters are supported:

- **Src Network Adrs**
The source IPX network address. Either the source or destination address (or both) must be specified.
- **Dst Network Adrs**
The destination IPX network address. Either the source or destination address (or both) must be specified.
- **Src Node Adrs**
A valid IPX node address. The node address `fffffffffff` means all nodes in the specified source network. This value must be specified if the Src Network Adrs is not null.
- **Dst Node Adrs**
A valid IPX node address. The node address `fffffffffff` means all nodes in the specified destination network. This value must be specified if the Dst Network Adrs is not null.
- **Src Socket Cmp and Src Socket #**
Some NetWare services communicate across specific sockets; for example, file servers typically use socket 0451. If you specify the source socket number, you can also specify the type of comparison to be made between the source socket for an IPX packet and the value specified in this filter. You can specify that the filter matches the packet if the source socket number is equal, not-equal, less-than, or greater-than the one specified in the filter.

- Dst Socket Cmp and Dst Socket #

If you specify the destination socket number, you can also specify the type of comparison to be made between the destination socket for an IPX packet and the value specified in this filter. You can specify that the filter matches the packet if the destination socket number is equal, not-equal, less-than, or greater-than the one specified in the filter.

Applying an IPX SAP filter

You can apply an IPX SAP filter to the local Ethernet or to WAN interfaces, or both.

- On Ethernet, a SAP filter includes or excludes specific servers or services from the table.

Open Ethernet > Mod Config > Ether Options.

If directory services is not supported, servers or services that are not in the Pipeline table will be inaccessible to clients across the WAN.

- In the Answer profile, a SAP filter screens service advertisements from across the WAN.

Open Ethernet > Answer > Session Options.

- In a Connection profile, a SAP filter screens service advertisements to and from a specific WAN connection.

Open Ethernet > Connections > *any profile* > Sessions Options.

To apply an IPX SAP filter profile:

- 1 Open the profile.
- 2 Open the Session Options submenu (Answer and Connection profiles) or Ether Options submenu (Ethernet profile).
- 3 Specify the number of the IPX SAP filter profile you defined.

You apply an IPX SAP Filter profile by specifying the unique part of the number it is assigned in the IPX SAP Filters menu. For example, to apply the filter defined as 20-801:

```
IPX SAP Filter=1
```

- 4 Close the profile.

A filter applied to the Ethernet interface takes effect immediately.

Configuring IPX routing connections

This section describes how to configure IPX routing connections. It describes typical host software requirements and includes the following example configurations:

- Example dial-in client connection
- Example with servers on both sides of the link
- Example with servers on only one side of the link

An example dial-in client connection

In Figure 4-1 a NetWare client dials into a corporate IPX network that supports both servers and clients using PPP dial-in software.




Figure 4-1. A dial-in client requiring dynamic IPX network assignment

In this example, the Pipeline is connected to a corporate NetWare LAN and the dial-in client has an ISDN modem, NetWare client software, and PPP dial-up software. This example assumes that the Answer profile and Ethernet profile have been set up to enable IPX routing.

To configure the Pipeline to accept a connection from the dial-in user:

- 1 Open the Ethernet profile.
- 2 Specify an IPX number for assignment to dial-in clients.

```
IPX Pool#=B21CC345
```

Note: Make sure this number is unique in the entire IPX routing domain.

- 3 Close the Ethernet profile.
- 4 Open the Connection profile for the dial-in user and set the following parameters.

```
Station=NetWareClient1
Active=Yes
Encaps=PPP
Route IPX=Yes

Encaps options...
    Send Auth=CHAP
    Recv PW=*SECURE*
    Send PW=*SECURE*

IPX options...
    Peer=Dialin
```

- 5 Close the Connection profile.

An example with NetWare servers on both sides of the link

In the following example the Pipeline is connected to an IPX network that supports both servers and clients. The example shows how it will make the connection to a remote site that also supports both servers and clients.




Figure 4-2. A connection with NetWare servers on both sides

In this example, site A and site B are both existing Novell LANs that implement NetWare 3.12 and NetWare 4 servers, NetWare clients, and a Pipeline. The NetWare server at site A is configured with the following information:

```
Name=SERVER-1
internal net CFC12345
```

Configuring IPX Routing

Configuring IPX routing connections

```
Load 3c509 name=ipx-card frame=ETHERNET_8023
Bind ipx ipx-card net=1234ABCD
```

The NetWare server at site B is configured as follows:

```
Name=SERVER-2
internal net 013DE888
Load 3c509 name=net-card frame=ETHERNET_8023
Bind ipx net-card net=9999ABFF
```

To configure the Pipeline at site A:

- 1 Assign the Pipeline a name if it does not already have one.
To assign the Pipeline a name, open the System profile and specify the name in the Name parameter. For example:

```
Name=SITEAGW
```

- 2 Open the Connection profile for site B.
For sake of example, the Connection profile for site B is profile #5. A profile's number is the unique part of the number it is assigned in the Connections menu. For example, the Connection profile defined as 20-105 is #5.

Set up the Connection profile as follows:

```
Station=SITEBGW
Active=Yes
Encaps=MPP
Dial #=555-1212
Route IP=No
Route IPX=Yes
Bridge=No
Dial brdcast=N/A

Encaps options...
    Send Auth=CHAP
    Recv PW=*SECURE*
    Send PW=*SECURE*

IPX options...
    IPX RIP=None
    IPX SAP=Both
    NetWare t/o=30
```


- 3 Close Connection profile #5.
- 4 Open the Ethernet profile and make sure that it is set up for IPX routing.

For example:

```
IPX Routing=Yes  
  
Ether options...  
    IPX Frame=802.2  
    IPX Enet #=1234ABCD
```

- 5 Close the Ethernet profile.

Because IPX RIP is set to None in the Connection profile, configure a static route to the remote server:

- 6 Open an IPX Route profile.
- 7 Set up a route to the remote NetWare server with the following settings:

```
Server Name=SERVER-2  
Active=Yes  
Network=013DE888  
Node=000000000001  
Socket=0451  
Server Type=0004  
Connection #=5
```

Note: The Connection # parameter in the IPX Route profile must match the number of the Connection profile you configured for connection to that site.

- 8 Close the IPX Route profile.

To configure the Pipeline at site B:

- 1 Assign the Pipeline a name if it does not already have one.
To assign the Pipeline a name, open the System profile and specify the name in the Name parameter. For example:

```
Name=SITEBGW
```

- 2 Open the Connection profile for site A.
For sake of example, the Connection profile for site A is profile #2. A profile's number is the unique part of the number it is assigned in the Connections menu. For example, the Connection profile defined as 20-102 is #2.

Configuring IPX Routing

Configuring IPX routing connections

Set up the Connection profile as follows:

```
Station=SITEAGW
Active=Yes
Encaps=MPP
Dial #=555-1213
Route IP=No
Route IPX=Yes
Bridge=No
Dial brdcast=N/A

Encaps options...
    Send Auth=CHAP
    Recv PW=*SECURE*
    Send PW=*SECURE*

IPX options...
    IPX RIP=None
    IPX SAP=Both
    NetWare t/o=30
```

- 3 Close Connection profile #2.
- 4 Open the Ethernet profile and make sure that it is set up for IPX routing.
For example:

```
IPX Routing=Yes

Ether options...
    IPX Frame=802.2
    IPX Enet #=9999ABFF
```

- 5 Close the Ethernet profile.

Because IPX RIP is set to None in the Connection profile, configure a static route to the remote server:

- 6 Open an IPX Route profile.
- 7 Set up a route to the remote NetWare server using these settings:

```
Server Name=SERVER-1
Active=Yes
Network=CFC12345
Node=000000000001
Socket=0451
```

```
Server Type=0004  
Connection #=2
```

Note: The Connection # parameter in the IPX Route profile must match the number of the Connection profile you configured to that site.

- 8 Close the IPX Route profile.

An example with local NetWare servers only

In the following example, the Pipeline is connected to a local IPX network that has both servers and clients, and the Pipeline will connect to a geographically remote network that supports one or more NetWare clients. Figure 4-3 shows the example setup.




Figure 4-3. A dial-in client that belongs to its own IPX network

In this example, site A implements NetWare 3.12 servers, NetWare clients, and a Pipeline. The NetWare server at site A is configured with the following information:

```
Name=SERVER-1  
internal net CFC12345  
Load 3c509 name=ipx-card frame=ETHERNET_8023  
Bind ipx ipx-card net=1234ABCD
```

Configuring IPX Routing

Configuring IPX routing connections

Site B is a home office that consists of one PC and a Pipeline. It is not an existing Novell LAN, so the Pipeline configuration creates a new IPX network (for example, 1000CFFF).

Note: The new IPX network number assigned to site B cannot be in use *anywhere* on the entire IPX wide-area network. (It cannot be in use at site A or any network to which site A connects.)

The example assumes that the Ethernet profile and Answer profile have already been set up to enable IPX routing. Because no static routes are used, the initial connection between the two Ascend units should be manually dialed (using the DO menu).

To configure the Pipeline at site A:

- 1 Assign the Pipeline a name if it does not already have one.
To assign the Pipeline a name, open the System profile and specify the name in the Name parameter. For example:

```
Name=SITEAGW
```

- 2 Open the Connection profile for site B.
Set up the Connection profile as follows:

```
Station=SITEBGW
Active=Yes
Encaps=MPP
Dial #=555-1212
Route IP=No
Route IPX=Yes
Bridge=No
Dial brdcast=N/A

Encaps options...
  Send Auth=CHAP
  Recv PW=*SECURE*
  Send PW=*SECURE*

IPX options...
  IPX RIP=Both
  IPX SAP=Both
  NetWare t/o=30
```

- 3 Close the Connection profile.

To configure the site B Ascend unit:

- 1** Assign the Ascend unit a name if it does not already have one.

To assign the Pipeline a name, open the System profile and specify the name in the Name parameter. For example:

```
Name=SITEBGW
```

- 2** Open the Connection profile for site A.

Set up the Connection profile like this:

```
Station=SITEAGW
```

```
Active=Yes
```

```
Encaps=MPP
```

```
Dial #=555-1213
```

```
Route IP=No
```

```
Route IPX=Yes
```

```
Bridge=No
```

```
Dial brdcast=N/A
```

```
Encaps options...
```

```
    Send Auth=CHAP
```

```
    Recv PW=*SECURE*
```

```
    Send PW=*SECURE*
```

```
IPX options...
```

```
    IPX RIP=Both
```

```
    IPX SAP=Both
```

```
    NetWare t/o=30
```

- 3** Close and save the profile.

Configuring the Pipeline as a Bridge

5

This chapter contains the following sections:

Introduction to Ascend bridging.	5-1
Enabling bridging.	5-6
Managing the bridge table	5-7
Configuring bridged connections	5-9

Introduction to Ascend bridging

In the Pipeline, bridges are used primarily to provide connectivity for protocols other than IP and IPX (AppleTalk, for example). They can also be used to join segments of an IP or IPX network. Because a bridging connection forwards packets at the hardware address level (link layer), it does not distinguish between protocol types and it requires no protocol-specific network configuration.

Bridging is very easy to configure and is commonly used to:

- Provide non-routed protocol connectivity with another site
- Link two sites so that their nodes appear to be on the same LAN
- Support protocols that depend on broadcasts to function, such as BOOTP

Be aware that bridges examine *all* packets on the LAN (called “promiscuous mode”), so they incur greater processor and memory overhead than routers. On heavily loaded networks, this increased overhead can result in slower performance.

Routing is much faster than bridging, and has these advantages:

- Routers examine packets at the network layer, so you can filter on logical addresses, providing enhanced security and control.
- Routers support multiple transmission paths to a given destination, enhancing the reliability and performance of packet delivery.

From a practical point of view, you should always route if possible, as routing is more efficient and makes call management easier. Bridging is necessary when you cannot subnet your IP network, and when you need to use non-routable protocols such as AppleTalk, NetBIOS, or DECnet.

How a bridged WAN connection is initiated

When the Pipeline is configured for bridging, it accepts all packets on the Ethernet and forwards only those that have one of the following:

- A physical address that is not on the segment connected to the Pipeline
- A broadcast address

Bridging uses physical or broadcast addresses, not logical (network) addresses.

Physical addresses and the bridge table

A physical address is a unique hardware-level address associated with a specific network controller. A device's physical address is also called its Media Access Control (MAC) address. On Ethernet, the physical address is a six-byte hexadecimal number assigned by the Ethernet hardware manufacturer, as in:

0000D801CFF2

If the Pipeline receives a packet whose destination MAC address is not on the local network, it checks its internal bridge table. If it finds the packet's MAC address, the Pipeline dials the connection and bridges the packet. If the address is *not* found, the Pipeline checks for active sessions that have bridging enabled. If there are active bridging links, the Pipeline forwards the packet across *all* active sessions that have bridging enabled.

Note: The Pipeline cannot dial a connection for packets that are not on the local network and not specified in its bridge table because it has no way of finding the proper Connection profile. See “Managing the bridge table” on page 5-7.

Broadcast addresses and Dial Brdcast

A broadcast address is recognized by multiple nodes on a network. For example, the Ethernet broadcast address at the physical level is:

FFFFFFFFFFFF

All devices on the same network receive all packets with that destination address. As a router, the Pipeline discards broadcast packets. As a bridge, it forwards packets with the broadcast destination address across all active sessions that have bridging enabled, and initiates a session for all Connection profiles in which the Dial Brdcast parameter is set to Yes.

Note: ARP broadcast packets that contain an IP address in the bridge table are a special case. For details, see “Static bridge-table entries” on page 5-8.

How bridged connections are established

Figure 5-1 show how station names and passwords sync a bridging connection.




Figure 5-1. Negotiating a bridge connection (PPP encapsulation)

The system name assigned to the Pipeline in the Name parameter of the System Profile must be exactly the same device name specified in the Connection profile on the remote bridge (the match is case sensitive). Similarly, the name assigned

to the remote bridge must be exactly the same name specified in the Station parameter of that Connection profile.

Note: The most common cause of trouble when initially setting up a PPP bridging connection is that the names are not specified exactly. Check for case, dashes, spaces, underscores, and so forth.

Bridging in the Answer profile

Bridging must be enabled on both the answering and dialing side of a PPP, MP, or MP+ session link. Otherwise the link cannot bridge packets. In addition, password authentication is required for unique identification of devices. Unlike an IP routing configuration, where the Pipeline uniquely identifies the calling device by its IP address, a bridging configuration does not include a built-in way to identify incoming callers. For that reason, password authentication is required unless IP routing is configured in the same Connection profile. For details about PPP or MPP encapsulation, see Chapter 1, “Configuring WAN Connections.”

To set Answer profile parameters for a bridging connection:

- 1 Open the Ethernet > Answer > PPP options menu.
- 2 Turn on bridging. (The setting for Bridge is N/A until the Bridging parameter in the Ethernet profile is set.)
`Bridge=Yes`
- 3 Set Recv Auth to Either (PAP, CHAP, or MS-CHAP).
- 4 Exit the Answer profile.

About IPX bridging

IPX bridging has special requirements for facilitating NetWare client/server logins across the WAN and preventing IPX RIP and SAP broadcasts from keeping a bridged connection up indefinitely.

Like all options in the IPX Options submenu, the Handle IPX parameter is set to N/A if an IPX frame type is not specified in the Ethernet profile. Also, if Route IPX is set to Yes in the Connection profile, the Handle IPX parameter is set to N/A, but acts as if it is set to Server.

When there is no server support on the local network

If the local Ethernet supports NetWare clients only and no NetWare servers, the bridging connection should enable a local client to bring up the WAN connection by querying (broadcasting) for a NetWare server on a remote network. However, the connection should not stay up indefinitely because of RIP or SAP broadcasts.

To accomplish this, open Ethernet > Connections > *profile* > IPX options and set Handle IPX=Client.

When there is no server support on the remote network

If the local network supports NetWare servers (or a combination of clients and servers) and the remote network supports NetWare clients only, the bridging connection should enable the Pipeline to respond to NCP watchdog requests for remote clients, but to bring down inactive connections whenever possible.

To accomplish this, open Ethernet > Connections > *profile* > IPX options and specify a timeout value (for example, set NetWare t/o=30), and set the Handle IPX parameter to Server.

When there is server support on both networks

If NetWare servers are supported on both sides of the WAN connection, it is strongly recommended that you use an IPX routing configuration instead of bridging IPX. If you bridge IPX in that type of environment, client/server logins are lost when the Pipeline brings down an inactive WAN connection.

IPX routing and bridging on the same connection

When IPX routing is enabled for a connection, the Pipeline routes only one packet frame type across that connection. For example, if the IPX frame type is set to 802.3, only 802.3 packets are routed. If some NetWare servers on the local network use a different frame type, such as 802.2, those packets are bridged if bridging is enabled, or discarded if bridging is *not* enabled.

Examples

If IPX Frame=802.3, and Route IPX=Yes and Bridge=No in the Connection profile, only 802.3 IPX packets are routed; all other packets are dropped.

If IPX Frame=802.3, and Route IPX=Yes and Bridge=Yes in the Connection profile, 802.3 IPX packets are routed and all other packets are bridged, including IPX packets in other frame types, AppleTalk packets, NetBios packets, DECnet and so forth.

If the Pipeline receives an IPX packet in the 802.2 packet frame, it uses the physical address in that packet to bridge it across all active bridging sessions.

Enabling bridging

The Pipeline has a global bridging parameter that must be enabled for any bridging connection to work. The Bridging parameter causes the Pipeline unit's Ethernet controller to run in promiscuous mode. In promiscuous mode, the Ethernet driver accepts all packets, regardless of address or packet type, and passes them up the protocol stack for a higher-layer decision on whether to route, bridge, or reject the packets.

Note: Running in promiscuous mode incurs greater processor and memory overhead than the standard mode of operation for the Ethernet controller. On heavily loaded networks, this increased overhead can result in slower performance, even if no packets are actually bridged.

To enable bridging on Ethernet:

- 1 Open the Ethernet > Mod Config > Ether Options.
- 2 Turn on the global bridging parameter.
Bridging=Yes
- 3 Close the Ethernet profile.

Managing the bridge table

To forward bridged packets to the right network destination, the Pipeline uses a bridge table that associates end nodes with particular connections. It builds this table dynamically, as described in “Transparent bridging” on page 5-7. It also incorporates the entries found in its Bridge profiles. Bridge profiles are analogous to static routes in a routing environment. You can define up to eight destination nodes and their connection information in Bridge profiles.

Parameters that affect the bridge table

Parameters directly related to the bridge table are set in the following menus:

```
Ethernet
  Mod Config
    Ether options...
      Bridging=Yes

Ethernet
  Connections
    profile
      Bridge=Yes
      Dial Brdcast=No

Ethernet
  Bridge Adrs
    Enet Adrs=CFD-12367
    Net Adrs=10.0.0.12
    Connection #=7
```

For details on each parameter, see the *Reference Guide*.

Transparent bridging

As a transparent (or learning) bridge, the Pipeline keeps track of where addresses are located as it forwards packets. It records each packet’s source address in a bridging table. A Connection profile is associated with an address when it is used to dial the link or when it matches an incoming call.

Figure 5-2 shows the physical addresses of some nodes on the local Ethernet and one at a remote site. The Pipeline at site A, configured as a bridge, gradually learns addresses on both networks by looking at each packet's source address.




Figure 5-2. How the Pipeline creates a bridging table

The resulting bridging table looks like this:

0000D801CFF2	SITEA
080045CFA123	SITEA
08002B25CC11	SITEA
08009FA2A3CA	SITEB (Connection Profile #5)

Entries in the Pipeline unit's bridge table must be relearned within a fixed aging time limit, or they are removed from the table.

Static bridge-table entries

The administrator can specify up to eight static bridge-table entries in Bridge profiles. Each connection that has a static bridge table entry can have the Dial Brdcast parameter set to No.

Dial Brdcast is a very convenient way of bridging packets if the Pipeline has only a few bridging connections, but it can be expensive in an environment where many profiles support bridging. (For more information, see "Broadcast addresses and Dial Brdcast" on page 5-3.) If Dial Brdcast is turned off in a Connection profile, the Pipeline does not initiate dialing for that connection on the basis of

broadcast requests. Instead, it relies on its bridging table to recognize which Connection profile to use.

Note: If you turn off Dial Brdcast and the Pipeline does not have a bridge-table entry for a destination address, the Pipeline will not bring up that connection.

To define a static bridge-table entry:

- 1 Open a Bridge profile.
- 2 Specify the physical address of the remote host.
For example:

```
Enet Adrs=0080AD12CF9B
```

Get this address from the administrator of the far-end device. For more information, see “Physical addresses and the bridge table” on page 5-2.
- 3 If the far-end is a segment of the local IP network, specify an address on that segment. For example:

```
Net Adrs=10.2.3.133
```

For more details, see “An example IP bridged connection” on page 5-16.
- 4 Specify the number of the Connection profile for this connection.
For example:

```
Connection #=2
```

You don’t have to specify the whole number, just the unique portion of it.
- 5 Exit and save the profile.

Configuring bridged connections

This section shows how to configure bridging for a Pipeline connecting to a remote site. The example configuration focuses on bridging. It does not show the link-specific settings (such as Telco options, MP+, or frame relay configuration), or additional routing settings that might be appropriate at your site.

Connection profiles must enable bridging, and if the remote network is not recorded as a static bridge-table entry, Dial Brdcast must also be enabled.

Parameters related to protocol-independent bridging are set in the following menus:

Configuring the Pipeline as a Bridge

Configuring bridged connections

```
Ethernet
Connections
  profile
    Station=SITEBGW
    Bridge=Yes
    Dial Brdcast=No
```

```
Ethernet
Connections
  profile
    Send Auth=None
    Recv PW=N/A
    Send PW=N/A
```

```
Ethernet
Connections
  profile
    IPX options...
    Handle IPX=Client
```

For details on each parameter, see the *Reference Guide*.

An example AppleTalk bridged connection

An AppleTalk connection at the link level requires a bridge at either end of the connection. Be careful when specifying names. Names are case sensitive, and dashes, spaces, underscores and other details must be retained. The most common cause of trouble when initially setting up a bridging connection is that the wrong name is specified for the Pipeline or the remote device. Make sure you type the name exactly as it appears in the remote device.

The following example assumes that Bridging has been enabled on the Ethernet interface (as discussed in “Enabling bridging” on page 5-6). It also assumes that

the Answer profile enables bridging (as discussed in “Bridging in the Answer profile” on page 5-4).

Note: In the example, Dial Brdcast is turned off in the Connection profiles and a Bridge profile is specified. This is not required. You can turn on Dial Brdcast and omit the Bridge profile if you prefer.

To configure the local Pipeline for a bridged connection:

- 1 Open the System profile.
- 2 If the Pipeline does not already have a system name, assign one.
For example:
Name=SITEAGW
Bridged connections use system names for part of the authentication process.
- 3 Close the System profile.
- 4 Open Connection profile #5.
- 5 Set these parameters:
Station=SITEBGW
Active=Yes
Encaps=PPP
Bridge=Yes
Dial Brdcast=No
Encaps options...
Send Auth=CHAP
Recv PW=*SECURE*
Send PW=*SECURE*
- 6 Close Connection profile #5.
- 7 Open a Bridge profile.
- 8 Set these parameters:
Enet Adrs=0080AD12CF9B
Net Adrs=0.0.0.0
Connection #=5
- 9 Close the Bridge profile.

To configure the remote Pipeline unit for a bridged connection:

- 1 Open the System profile (on the remote Pipeline).

Configuring the Pipeline as a Bridge

Configuring bridged connections

- 2 If the Pipeline does not already have a system name, assign one.

For example:

```
Name=SITEBGW
```

- 3 Close the System profile.
- 4 Open Connection profile #2 on the Pipeline.
- 5 Set these parameters:

```
Station=SITEAGW
```

```
Active=Yes
```

```
Encaps=PPP
```

```
Bridge=Yes
```

```
Dial Brdcast=No
```

```
Encaps option...
```

```
    Send Auth=CHAP
```

```
    Recv PW=*SECURE*
```

```
    Send PW=*SECURE*
```

- 6 Close Connection profile #2.
- 7 Open a Bridge profile.
- 8 Set these parameters:

```
Enet Adrs=0CFF1238FFFF
Net Adrs=0.0.0.0
Connection #=2
```
- 9 Close the Bridge profile.

An example IPX client bridge (local clients)

In the following example, the local Ethernet supports NetWare clients, and the remote network supports NetWare servers and clients.




Figure 5-3. An example IPX client bridging connection

To configure the Pipeline in this example:

- 1 Open the System profile.
- 2 If the Pipeline does not already have a system name, assign one.

For example:

Name=SITEAGW

- 3 Close the System profile.
- 4 Open the Ethernet profile.
- 5 Open the Ether Options submenu.
- 6 Set the IPX Frame type.

IPX Frame=802.3

- 7 Close the Ethernet profile.
- 8 Open a Connection profile.

- 9 Set these parameters:

Station=SITEBGW

Active=Yes

Encaps=PPP

Route IPX=No

Bridge=Yes

Dial Brdcast=Yes

Configuring the Pipeline as a Bridge

Configuring bridged connections

```
Encaps options...
  Send Auth=CHAP
  Recv PW=*SECURE*
  Send PW=*SECURE*

IPX options...
  Handle IPX=Client
```

- 10 Close the Connection profile.

Dial Brdcast is enabled to allow service queries to bring up the connection.

When Handle IPX=Client, the Pipeline applies a data filter that discards RIP and SAP periodic broadcasts at its WAN interface, but forwards RIP and SAP queries. That way, local clients can locate a NetWare server across the WAN, but routine broadcasts do not keep the connection up unnecessarily.

An example IPX server bridge (local servers)

In the following example, the local network supports a combination of NetWare clients and servers, and the remote network only supports clients.




Figure 5-4. An example IPX server bridging connection

To configure the Pipeline in this example:

- 1 Open the System profile.
- 2 If the Pipeline does not already have a system name, assign one.
For example:
Name=SITEAGW
- 3 Close the System profile.

- 4 Open the Ethernet profile.
- 5 Open the Ether Options submenu.
- 6 Set the IPX Frame type.

For example:

```
IPX Frame=802.3
```

- 7 Close the Ethernet profile.
- 8 Open a Connection profile.
- 9 Set these parameters:

```
Station=SITEBGW
```

```
Active=Yes
```

```
Encaps=PPP
```

```
Route IPX=No
```

```
Bridge=Yes
```

```
Dial Brdcast=Yes
```

```
Encaps options...
```

```
Send Auth=CHAP
```

```
Recv PW=*SECURE*
```

```
Send PW=*SECURE*
```

```
IPX options...
```

```
NetWare t/o=30
```

```
Handle IPX=Server
```

- 10 Close the Connection profile.

When Handle IPX=Server, the Pipeline applies a data filter that discards RIP and SAP broadcasts at its WAN interface, but forwards RIP and SAP queries. It also uses the value specified in the “NetWare t/o” parameter as the time limit for responding to NCP watchdog requests on behalf of clients on the other side of the bridge, a process called “watchdog spoofing.”

Note: The Pipeline performs watchdog spoofing for the IPX frame type specified in the Ethernet profile. For example, if IPX Frame=802.3, only connections to servers using that packet frame type will be spoofed. (For more information, see Chapter 4, “Configuring IPX Routing.”)

An example IP bridged connection

If you are bridging between two segments of the same IP network, you can use the Net Adrs parameter in a Bridge profile to enable the Pipeline to respond to ARP requests while bringing up the bridged connection.

If an ARP packet contains an IP address that matches the Net Adrs parameter of a Bridge profile, the Pipeline responds to the ARP request with the Ethernet (physical) address specified in the Bridge profile, and brings up the specified connection. In effect, the Pipeline acts as a proxy for the node that actually has that address.

In this example, two segments of an IP network are connected across the WAN.




Figure 5-5. An example IP bridging connection

To configure the first Pipeline shown in Figure 5-5:

- 1 Open the System profile.
- 2 If the Pipeline does not already have a system name, assign one.
For example:

Name=SITEAGW

- 3 Close the System profile.
- 4 Open Connection profile #7 (for example).
- 5 Set these parameters:

Station=SITEBGW

Active=Yes

Encaps=PPP

Route IP=No

Bridge=Yes

Dial Brdcast=No

```
Encaps options...  
  Send Auth=CHAP  
  Recv PW=*SECURE*  
  Send PW=*SECURE*
```

6 Close Connection profile #7.

7 Open a Bridge profile.

8 Set these parameters:

```
Enet Adrs=0CFF1238FFFF  
Net Adrs=10.2.3.100/24  
Connection #=7
```

9 Close the Bridge profile.

Defining Filters and Firewalls

This chapter contains the following topics:

Introduction to filters	6-1
Overview of Filter profiles	6-6
Example filters	6-12
Working with predefined call filters	6-21
Display unwanted dial-out packets.	6-28
Secure Access Firewalls	6-34
Filter persistence.	6-36

Introduction to filters

Filters inspect packets, and depending on the attributes of the packet, filters reject packets from entering or leaving your network. When a filter is in use, the Pipeline examines every packet in the packet stream and takes action if the defined filter conditions are present. The action the Pipeline takes depends both on the conditions specified within the filter and how the filter is applied.

The default action when no filter is used is to forward (accept) all packets and allow all packets to reset the idle timer, which is used to determine when to disconnect inactive sessions.

You can define conditions in filters to drop (reject) all packets except the ones you explicitly allow, or allow all packets except the ones you explicitly drop. Additionally, you can specify whether to apply the filter to inbound packets, outbound packets, or all packets, regardless of their origin.

Depending on how a filter is used, it is either a data filter or a call filter. The following describes each type:

- **Data filter**
Affects the flow of data. Packets are dropped (rejected) or forwarded (accepted) as specified in the filter conditions. Mainly used for security.
- **Call filter**
Determines which packets can initiate a connection or reset the idle timer for an established connection. Mainly used to prevent unnecessary connections.

Note: Packets can pass through more than one filter. If both a data filter and call filter are applied, the data filter takes precedence.

Data filters for dropping or forwarding certain packets

Data filters are commonly used for security, but they can be used for any purpose that requires the Pipeline to drop or forward specific packets. For example, you can use data filters to drop packets addressed to particular hosts, or to prevent broadcasts from going across the WAN. You can also use data filters to allow only specified devices to be accessed by users across the WAN.

Data filters do not affect the idle timer, and a data filter applied to a Connection profile does not affect the answering process.




Figure 6-1. Data filters can drop or forward certain packets

To define which packets will be allowed to cross the WAN interface, apply a data filter to a Connection or Answer profile using the following steps:

- 1 Open Ethernet > Connection or Answers > *profile*

Note: You can apply a filter in the Answer profile only if the Profile Req'd parameter is set to No.

- 2 Open the Session Options submenu.

- 3 Apply a data filter.

For example:

Data Filter=4

If this parameter is set to zero, the default, no filter is applied. To apply a filter, specify its profile number. You can view the profile number by opening the Filters menu. You don't have to specify the whole number, just the unique portion of it, for example, 1, 2, 3,...

- 4 Close and save the profile.

A filter applied to a Connection profile takes effect only when the connection goes from an offline state to a call-placed state.

To define which packets will be allowed to cross the Ethernet interface, apply a data filter to a Connection profile using the following steps:

- 1 Open Ethernet > Mod Config > Ether Options.

- 2 Apply the data filter.

For example:

Data Filter=4

If this parameter is set to zero, the default, no filter is applied. To apply a filter, specify its profile number. You can view the profile number by opening the Filters menu. You don't have to specify the whole number, just the unique portion of it, for example, 1, 2, 3,...

- 3 Close and save the profile.

A filter applied to the Ethernet interface takes effect immediately. If you change any of the conditions in the Filter profile definition, new or changed conditions are applied as soon as you save the Filter profile.

For an example data filter, see "Example filters" on page 6-12.

Call filters for managing connections

Call filters are used to prevent unnecessary connections and to help the Pipeline distinguish active traffic from “noise.” By default, any traffic to a remote site triggers a call to that site, and any traffic across an active connection resets the connection’s idle timer.

Note: The idle timer is set to 120 seconds by default. If a connection is inactive for two minutes, the idle timer expires and the Pipeline terminates the connection.

Call filters define which packets are not considered active traffic on a particular connection. They identify which packets should not originate a connection or reset the idle timer. Call filters do not affect which packets are transmitted or received across active connections.




Figure 6-2. Call filters used to prevent resetting the timer

To define which packets will not reset the idle timer or keep a connection active, apply a call filter to a Connection or Answer profile using the following steps:

- 1 Open Ethernet > Connection or Answer > *profile*

Note: You can apply a filter in the Answer profile only if the Profile Req'd parameter is set to No.

- 2 Open the Session Options submenu.
- 3 Apply the call filter.

For example:

Call Filter=5

If this parameter is set to zero, the default, no filter is applied.

If it is set to any other value, the value must be a valid Filter profile number. The Filter profile number is the number in the Filters menu. You don't have to specify the whole number, just the unique portion of it.

- 4 Close and save the profile.

When you apply a filter to the WAN interface, it takes effect only when a connection goes from an offline state to a call-placed state.

To reset the idle timer, perform the following steps:

- 1 Open Ethernet > Connection or Answer > *profile*

Note: You can apply a filter in the Answer profile only if the Profile Req'd parameter is set to No.

- 2 Open the Session Options submenu.
- 3 Specify the number of seconds to wait before clearing an inactive connection.

For example:

`Idle=15`

If this parameter is set to zero, an idle connection stays open indefinitely. For example, if you specify 15, an idle connection is terminated after 15 seconds.

- 4 Close and save the profile.

Predefined call filters

The Pipeline ships with the following predefined Filter profiles:

- IP Call, for IP connections.
- NetWare Call, for IPX connections.
- AppleTalk Call, for bridged AppleTalk connections.

These filters are basic call filters that prevent the most common traffic in each kind of packet stream from initiating or maintaining a connection. (For

information about predefined-filter settings, see “Working with predefined call filters” on page 6-21.)

Note: For information about IPX SAP filters, pertaining to NetWare services the Pipeline adds to its service table, see Chapter 4, “Configuring IPX Routing.”

Overview of Filter profiles

You apply a filter to an interface by specifying its profile number. The Pipeline applies all filter conditions defined in a Filter profile to the Connection or Answer profile where it is specified.

Figure 6-3 shows how filters are organized in the menu interface, and the terminology used to describe each part of a filter.




Figure 6-3. Filter organization and terminology

The menus shown in Figure 6-3 are nested, beginning with the Filters menu. That is, the numbered Filter profile menus are found under the Filters menu; the Input or Output filter menus are found under each numbered Filter profile menu, and so forth. Each level of the Filters menu is described as follows:

- **Filters menu**
The Filters menu contains a list of numbered profiles. When applying a filter, you identify it by the unique portion of its Filter profile number, for example, you would use 1, 2, or 3, rather than 20-401, 20-402, or 20-403.
- **Filter profile**
A Filter profile is a set of filter conditions.
- **Input or Output filters**

At the top level of a Filter profile are two submenus: Input Filters and Output Filters. The Input submenu allows you to define 12 In-filter conditions to apply to incoming data. The Output submenu allows you to define 12 Out-filter conditions to apply to outgoing data. The conditions are applied to the data stream in filter order, starting with 01.

- Generic or IP filters

Each In filter and Out filter can be one of two types: “Generic” or “IP.” After assigning a type, you define filter conditions applicable to that type of packet in its corresponding submenu.

- Filter conditions

Filter conditions specify the actual packet characteristics that will be examined in the data stream. Generic filter conditions specify locations and values that may be found within any packet. IP filter conditions specify packet characteristics that apply only to TCP/IP/UDP packets, such as address, mask, and port.

Filtering inbound and outbound packets

At the top level of a Filter profile, you can assign a name and open the Input Filters or Output Filters submenu.

```
20-401 IP Call
>Name=IP Call
  Input filters...
  Output filters...
```

Input filters cause the Pipeline to examine incoming packets, and Output filters cause it to examine outgoing packets. If the filter is applied as a data filter on the Ethernet, it affects packets from the Ethernet *into* the Pipeline or from the Pipeline *out* to the Ethernet. If applied as a data or call filter on a WAN interface defined in a Connection or Answer profile, it affects packets from that WAN interface *into* the Pipeline or from the Pipeline *out* to that interface.

You can specify up to 12 In filters and 12 Out filters in a Filter profile. These filters are applied in filter-number order, beginning with In filter 01.

```
20-401 IP Call
  Input filters...
    >In filter 01
```

```
In filter 02
In filter 03
In filter 04
In filter 05
In filter 06
In filter 07
In filter 08
In filter 09
In filter 10
In filter 11
In filter 12
```

By default, all packets are forwarded. So if a packet does not match any of the defined conditions in a filter, it is forwarded as usual.

Note: If only Input filters are defined, all outbound packets are forwarded or allowed to reset the idle timer. If you define only Output filters, all inbound packets are forwarded or allowed to reset the idle timer.

Selecting filter type and activating the filter

The In filters and Out filters you define are applied to a packet in the order in which they appear in this list, provided that each filter has the Valid parameter set to Yes. Setting the Valid parameter to No prevents it from being applied.

When you open an “In filter” or an “Out filter,” set the Valid parameter to Yes and select the type of filter conditions to be defined, Generic, IP, or IPX.

```
20-401 IP Call
In filter 01
>Valid=Yes
Type=GENERIC
Generic...
IP...
IPX...
```

Generic filter conditions define bits and bytes within a packet. They are applied to all packet types, including TCP and IP. IP filter conditions are related only to TCP/IP/UDP packets.

Defining generic filter conditions

If the Type parameter in a filter is set to GENERIC, you can define generic conditions using these menus:

- 1 Open Ethernet > Filters > *any profile*
- 2 Determine if you need an Input or Output filter.
- 3 Open a filter, from 01 to 12, and select Generic

For example:

```
Forward=No
Offset=14
Length=8
Mask=ffffffffffffffff
Value=aaa03000000080f3
Compare=Equals
More=No
```

- 4 Set the Forward parameter. It determines whether the Pipeline will forward a packet if it matches the definition, where Forward=Yes, or drop the packet if it matches, where Forward=No.

If a filter is applied as a data filter, the “forward” action determines which packets will be transmitted and received. If a filter is applied as a call filter, the “forward” action determines which packets can either initiate a connection or reset the timer for an established connection.

- 5 Set Offset, Length, Mask, and Value.

The Offset, Length, Mask, and Value parameters are used to define the exact location of certain bytes within a packet and the value of those bytes.

- 6 Set the Compare parameter.

The Compare parameter specifies how a packet’s contents are compared to the value specified in this filter. After applying the Offset, Mask, and Length values to reach the appropriate location in a packet, the contents of that location are compared to the Value parameter. If Compare is set to Equals, the default, the filter is applied if the packet data are identical to the specified

value. If Compare is set to NotEquals, the filter is applied if the packet data are not identical.

7 Set the More parameter.

The More parameter specifies whether the current filter is linked to the one immediately following it. If More=Yes, the filter can examine multiple non-contiguous bytes within a packet, by “marrying” the current filter to the next one, so that the next filter is applied before the Forward decision is made. The match occurs only if *both* non-contiguous bytes contain the specified values. If More=No, the Forward decision is based on whether the packet matches the definition in this one filter.

Defining IP filter conditions

If the Type parameter is set to IP, you can define filter conditions related only to TCP/IP/UDP data packets, including bridged packets using these menus:

- 1** Open Ethernet > Filters > *any profile*
- 2** Determine if you need an Input or Output filter.
- 3** Open a filter, from 01 to 12, and select IP

For example:

```
Forward=Yes
Src Mask=255.255.255.192
Src Adrs=192.100.40.128
Dst Mask=0.0.0.0
Dst Adrs=0.0.0.0
Protocol=0
Src Port Cmp=None
Src Port #=N/A
Dst Port Cmp=None
Dst Port #=N/A
TCP Estab=N/A
```

An IP filter examines source addresses, destination addresses, IP protocol type and port, or a combination of these.

4 Set the Forward parameter.

The Forward parameter determines whether the Pipeline will forward a packet if it matches the definition, where Forward=Yes, or drop the packet if it matches, where Forward=No.

If a filter is applied as a data filter, the “forward” action determines which packets will be transmitted and received. If a filter is applied as a call filter, the “forward” action determines which packets can either initiate a connection or reset the timer for an established connection.

5 Set the source and destination address and mask/

The source and destination Mask and Adrs parameters specify the contents of the source or destination fields in a packet. Use the Mask parameter to mask out portions of the source or destination address, for example, to mask out the host number.

6 Specify the Protocol.

The Protocol parameter is used to identify a specific TCP/IP protocol; for example, 6 specifies TCP packets. Common protocols are listed below, but protocol numbers are not limited to this list. For a complete list, see the section on Well-Known Port Numbers in RFC 1700, *Assigned Numbers*, by Reynolds, J. and Postel, J., October 1994.

- 1 — ICMP
- 5 — STREAM
- 8 — EGP
- 6 — TCP
- 9 — Any private interior gateway protocol, such as Cisco’s IGRP
- 11 — Network Voice Protocol
- 17 — UDP
- 20 — Host Monitoring Protocol
- 22 — XNS IDP
- 27 — Reliable Data Protocol
- 28 — Internet Reliable Transport Protocol
- 29 — ISO Transport Protocol Class 4
- 30 — Bulk Data Transfer Protocol
- 61 — Any Host Internal Protocol
- 89 — OSPF

7 Set the source and destination ports and comparison method.

The source and destination Port Cmp and Port # parameters specify whether to compare the protocol ports, which identify the application running over TCP/IP. The comparison may match a protocol port number that is less-than, greater-than, equal, or not-equal.

8 Set the TCP Estab parameter.

The TCP Estab parameter can be set to match a packet only if a TCP session is already established.

Example filters

This section provides a step-by-step examples of defining filters. It shows how to specify both generic and IP filter conditions.

This section shows how to create Filter profiles. Some sites modify the predefined call filters to make them more full-featured for the types of packets commonly seen at that site. See “Working with predefined call filters” on page 6-21 for details.

An example generic filter to handle AppleTalk broadcasts

This section shows how to define a generic data filter whose purpose is to prevent local AppleTalk AEP and NBP traffic from going across the WAN. The data filter first defines the types of packets that should *not* be filtered:

- AppleTalk Address Resolution Protocol (ARP) packets
- AppleTalk packets that are not addressed to the AppleTalk multicast address, such as regular traffic related to an actual AppleTalk File Server connection
- All non-AppleTalk traffic

The filter then defines the packets that should be dropped:

- AppleTalk Echo Protocol (AEP)
- Name Binding Protocol (NBP)

To define a generic data filter:

- 1** Select an unnamed Filter profile in the Filters menu and press Enter.
For example, select 20-403.

- 2 Assign a name to the Filter profile.

For example:

Name=AppleTalk Data

- 3 Open the Output Filters submenu.
- 4 Open Out filter 01 and set Valid=Yes and Type=GENERIC.

For example:

```
>Valid=Yes
  Type=GENERIC
  Generic...
  IP...
  IPX...
```

- 5 Open the Generic submenu and specify the following conditions:

```
Generic...
  >Forward=No
  Offset=14
  Length=8
  Mask=fffffffffffffffff
  Value=aaaa0300000080f3
  Compare=Equals
  More=No
```

These conditions define a location within a packet and the hexadecimal value that AARP packets contain within that location, protocol type 0x80f3. Outbound AARP packets will not be forwarded.

- 6 Close Out filter 01, and then open Out filter 02.
- 7 Set Valid=Yes and Type=GENERIC, and then open the Generic submenu and specify the following conditions:

```
Generic...
  >Forward=Yes
  Offset=14
  Length=8
  Mask=fffffffffffffffff
  Value=aaaa030800007809b
  Compare=NotEquals
  More=No
```

Defining Filters and Firewalls

Example filters

These conditions define non-AppleTalk traffic. Note that AppleTalk has the protocol type 0x809b. Outbound packets that are not AppleTalk packets will be forwarded. Because all non-AppleTalk packets have now been forwarded, subsequent filters can assume that a packet is AppleTalk.

- 8 Close Out filter 02, then open Out filter 03.
- 9 Set Valid=Yes and Type=GENERIC, and then open the Generic submenu and specify the following conditions:

```
Generic...
>Forward=Yes
  Offset=32
  Length=3
  Mask=ffffff0000000000
  Value=0404040000000000
  Compare=Equals
  More=No
```

These conditions filter AEP packets.

- 10 Close Out filter 03, then open Out filter 04.
- 11 Set Valid=Yes and Type=GENERIC, and then open the Generic submenu and specify the following conditions:

```
Generic...
>Forward=Yes
  Offset=32
  Length=6
  Mask=ffffffffffff0000
  Value=090007ffffff0000
  Compare=NotEquals
  More=No
```

AppleTalk “broadcast” traffic uses a multicast address. These conditions specify the multicast address. Any AppleTalk packet that does not use the multicast address will be forwarded.

- 12 Close Out filter 04, then open Out filter 05.
- 13 Set Valid=Yes and Type=GENERIC, and then open the Generic submenu and specify the following conditions:

```
Generic...
>Forward=Yes
  Offset=32
```

```

Length=4
Mask=ff00fff000000000
Value=0200022000000000
Compare=Equals
More=Yes

```

Together, Out filters 05 and 06 specify NBP lookup packets with a wildcard entity name. NBP lookups are transmitted by the Chooser and other applications that look up entities on AppleTalk networks.

- 14** Close Out filter 05, then open Out filter 06.
- 15** Set Valid=Yes and Type=GENERIC, and then open the Generic submenu and specify the following conditions:

```

Generic...
>Forward=Yes
Offset=42
Length=2
Mask=ffff000000000000
Value=013d000000000000
Compare=Equals
More=No

```

- 16** Close Out filter 06, then open Out filter 07.
- 17** Set Valid=Yes.

To discard everything else, just set Valid to Yes. This causes the default settings shown below:

```

Generic...
>Forward=No
Offset=0
Length=0
Mask=0000000000000000
Value=0000000000000000
Compare=Equals
More=No

```

- 18** Close and save the Filter profile.

An example IP filter to prevent address spoofing

This section shows how to define an IP data filter whose purpose is to prevent “spoofing” of local IP addresses. “Spoofing” IP addresses—not to be confused with watchdog or DHCP spoofing described elsewhere in this manual—is a technique whereby outside users pretend to be from the local network in order to obtain unauthorized access to the network.

The filter first defines Input filters that drop packets whose source address is on the local IP network or the loopback address (127.0.0.0). In effect, these filters say: “If you see an inbound packet with one of these source addresses, drop the packet.” The third Input filter defines every other source address (0.0.0.0) and specifies “Forward everything else to the local network.”

The data filter then defines an Output filter that specifies: “If an outbound packet has a source address on the local network, forward it; otherwise, drop it.” All outbound packets with a non-local source address will be dropped.

Note: This example assumes a local IP network address of 192.100.50.128, with a subnet mask of 255.255.255.192. Of course, you’ll use your own local IP address and netmask when defining a Filter profile.

Note: Because the Pipeline only supports 3 filters, this example modifies the predefined IP Call filter. See “Working with predefined call filters” on page 6-21 for information about predefined filters.

To define an IP data filter:

- 1 Select an unnamed Filter profile in the Filters menu and press Enter.

For example, select 20-401.

```
20-400 Filters
20-401 IP Call
20-402 NetWare Call
20-403 AppleTalk Call
```

- 2 Assign a name to the Filter profile.

For example:

```
Name=no spoofing
```

- 3 Open the Input Filters submenu.

4 Open In filter 01.

```
In filter 01
>Valid=Yes
Type=IP
Generic...
IP...
IPX...
```

5 Set Valid=Yes and Type=IP, and then open the IP submenu.

6 Specify the following conditions:

```
Ip...
>Forward=No
Src Mask=255.255.255.192
Src Adrs=192.100.50.128
Dst Mask=0.0.0.0
Dst Adrs=0.0.0.0
Protocol=0
Src Port Cmp=None
Src Port #=N/A
Dst Port Cmp=None
Dst Port #=N/A
TCP Estab=N/A
```

These conditions specify the local net mask and IP address in the Src Mask and Src Adrs fields. If an incoming packet has the local address, it will not be forwarded onto the Ethernet.

7 Close the current Input filter, and then open In filter 02.

8 Set Valid=Yes and Type=IP, and then open the IP submenu and specify the following conditions:

```
Ip...
>Forward=No
Src Mask=255.0.0.0
Src Adrs=127.0.0.0
Dst Mask=0.0.0.0
Dst Adrs=0.0.0.0
Protocol=0
Src Port Cmp=None
Src Port #=N/A
Dst Port Cmp=None
```

Defining Filters and Firewalls

Example filters

```
Dst Port #=N/A
```

```
TCP Estab=N/A
```

These conditions specify the loopback address in the Src Mask and Src Adrs fields. If an incoming packet has this address, it will not be forwarded onto the Ethernet.

- 9 Close the current Input filter, and then open In filter 03.
- 10 Set Valid=Yes and Type=IP, and then open the IP submenu and specify the following conditions:

```
Ip...
>Forward=Yes
Src Mask=0.0.0.0
Src Adrs=0.0.0.0
Dst Mask=0.0.0.0
Dst Adrs=0.0.0.0
Protocol=0
Src Port Cmp=None
Src Port #=N/A
Dst Port Cmp=None
Dst Port #=N/A
TCP Estab=N/A
```

These conditions specify every other source address (0.0.0.0) If an incoming packet has any non-local source address, it will not be forwarded onto the Ethernet.

- 11 Close the Input filter, and then return to the top level of the “no spoofing” Filter profile.
- 12 Open the Output Filters submenu, and select Out filter 01.
- 13 Set Valid=Yes and Type=IP, and then open the IP submenu and specify the following conditions:

```
Ip...
>Forward=Yes
Src Mask=255.255.255.192
Src Adrs=192.100.40.128
Dst Mask=0.0.0.0
Dst Adrs=0.0.0.0
Protocol=0
Src Port Cmp=None
Src Port #=N/A
```



```
Dst Port Cmp=None
Dst Port #=N/A
TCP Estab=N/A
```

These conditions specify the local net mask and IP address in the Src Mask and Src Adrs fields. If an outbound packet has a local source address, it will be forwarded.

14 Close the Filter profile.

An example IP filter for more complex security issues

This section describes an IP data filter that illustrates some of the issues you may need to consider when writing your own IP filters. The sample filter does not address fine points of network security. You may want to use this sample filter as a starting point and augment it to address your security requirements.

In this example, the local network supports a Web server and the administrator needs to provide dial-in access to the server's IP address while restricting dial-in traffic to all other hosts on the local network. However, many local IP hosts need to dial out to the Internet and use IP-based applications such as Telnet or FTP, which means that their response packets need to be directed appropriately to the originating host. In this example, the Web server's IP address is 192.9.250.5.

This filter would be applied as a data filter in a Connection or Answer profile.

```
In filter 01...Ip...Forward=Yes
In filter 01...Ip...Src Mask=0.0.0.0
In filter 01...Ip...Src Adrs=0.0.0.0
In filter 01...Ip...Dst Mask=255.255.255.255
In filter 01...Ip...Dst Adrs=192.9.250.5
In filter 01...Ip...Protocol=6
In filter 01...Ip...Src Port Cmp=None
In filter 01...Ip...Src Port #=N/A
In filter 01...Ip...Dst Port Cmp=Eq1
In filter 01...Ip...Dst Port #=80
In filter 01...Ip...TCP Estab=No

In filter 02...Ip...Forward=Yes
In filter 02...Ip...Src Mask=0.0.0.0
In filter 02...Ip...Src Adrs=0.0.0.0
```

Defining Filters and Firewalls

Example filters

```
In filter 02...Ip...Dst Mask=0.0.0.0
In filter 02...Ip...Dst Adrs=0.0.0.0
In filter 02...Ip...Protocol=6
In filter 02...Ip...Src Port Cmp=None
In filter 02...Ip...Src Port #=N/A
In filter 02...Ip...Dst Port Cmp=Gtr
In filter 02...Ip...Dst Port #=1023
In filter 02...Ip...TCP Estab=No
```

```
In filter 03...Ip...Forward=Yes
In filter 03...Ip...Src Mask=0.0.0.0
In filter 03...Ip...Src Adrs=0.0.0.0
In filter 03...Ip...Dst Mask=0.0.0.0
In filter 03...Ip...Dst Adrs=0.0.0.0
In filter 03...Ip...Protocol=17
In filter 03...Ip...Src Port Cmp=None
In filter 03...Ip...Src Port #=N/A
In filter 03...Ip...Dst Port Cmp=Gtr
In filter 03...Ip...Dst Port #=1023
In filter 03...Ip...TCP Estab=No
```

```
In filter 04...Ip...Forward=Yes
In filter 04...Ip...Src Mask=0.0.0.0
In filter 04...Ip...Src Adrs=0.0.0.0
In filter 04...Ip...Dst Mask=0.0.0.0
In filter 04...Ip...Dst Adrs=0.0.0.0
In filter 04...Ip...Protocol=1
In filter 04...Ip...Src Port Cmp=None
In filter 04...Ip...Src Port #=N/A
In filter 04...Ip...Dst Port Cmp=None
In filter 04...Ip...Dst Port #=N/A
In filter 04...Ip...TCP Estab=No
```

The first Input filter specifies the Web server's IP address as the destination and sets IP forward to Yes, so all IP packets received with that destination address will be forwarded.

The second Input filter specifies TCP packets, Protocol=6, *from* any address and *to* any address and forwards them if the destination port is greater than the source port. For example, Telnet requests go out on port 23 and responses come back on some random port greater than port 1023. So, this filter defines packets coming back to respond to a user's request to Telnet, or to other requests using the TCP protocol, to a remote host.

The third Input filter specifies UDP packets, Protocol=17, with exactly the same situation as described above for Telnet. For example, a RIP packet is sent out as a UDP packet to destination port 520. The response to this request also is sent to a random destination port greater than 1023.

Finally, the fourth Input filter specifies unrestricted pings and traceroutes. ICMP does not use ports like TCP and UDP, so a port comparison is unnecessary.

Working with predefined call filters

The Pipeline ships with three predefined Filter profiles, one for each commonly used protocol suite.

- IP Call, for IP connections
- NetWare Call, for IPX connections
- AppleTalk Call, for bridged AppleTalk connections

These predefined filters are intended as call filters, to help keep connectivity costs down. They provide a base that you can build on to fine-tune how the Pipeline handles routine traffic on your network.

Note: You can modify the predefined Filter profiles to make them more full-featured for the types of packets commonly seen on your network that you want to prevent from initiating or maintaining connections.

NetWare Call filter

The predefined NetWare Call filter is designed to prevent Service Advertising Protocol (SAP) packets originating on the local IPX network from resetting the idle timer or initiating a call.

Defining Filters and Firewalls

Working with predefined call filters

NetWare servers broadcast SAP packets every 60 seconds to make sure that all routers and bridges know about available services. To prevent these packets from keeping a connection up unnecessarily, apply the predefined NetWare Call filter in the Session Options submenu of a Connection or Answer profile in which IPX routing is configured.

The predefined NetWare Call filter contains six Output filters, which identify outbound SAP packets and prevent them from resetting the idle timer or initiating a call.

```
Out filter 01...Generic...Forward=No
Out filter 01...Generic...Offset=14
Out filter 01...Generic...Length=3
Out filter 01...Generic...Mask=ffffff000000000000
Out filter 01...Generic...Value=e0e0030000000000
Out filter 01...Generic...Compare=Equals
Out filter 01...Generic...More=Yes
```

```
Out filter 02...Generic...Forward=No
Out filter 02...Generic...Offset=27
Out filter 02...Generic...Length=8
Out filter 02...Generic...Mask=ffffffffffffff
Out filter 02...Generic...Value=ffffffffffff0452
Out filter 02...Generic...Compare=Equals
Out filter 02...Generic...More=Yes
```

```
Out filter 03...Generic...Forward=No
Out filter 03...Generic...Offset=47
Out filter 03...Generic...Length=2
Out filter 03...Generic...Mask=ffff000000000000
Out filter 03...Generic...Value=0002000000000000
Out filter 03...Generic...Compare=Equals
Out filter 03...Generic...More=No
```

```
Out filter 04...Generic...Forward=No
Out filter 04...Generic...Offset=12
Out filter 04...Generic...Length=4
Out filter 04...Generic...Mask=fc00ffff00000000
Out filter 04...Generic...Value=0000ffff00000000
```

```
Out filter 04...Generic...Compare=Equals
Out filter 04...Generic...More=Yes

Out filter 05...Generic...Forward=No
Out filter 05...Generic...Offset=24
Out filter 05...Generic...Length=8
Out filter 05...Generic...Mask=ffffffffffffffff
Out filter 05...Generic...Value=ffffffffffff0452
Out filter 05...Generic...Compare=Equals
Out filter 05...Generic...More=Yes

Out filter 06...Generic...Forward=No
Out filter 06...Generic...Offset=44
Out filter 06...Generic...Length=2
Out filter 06...Generic...Mask=ffff000000000000
Out filter 06...Generic...Value=0002000000000000
Out filter 06...Generic...Compare=Equals
Out filter 06...Generic...More=No
```

Extending the predefined filter for RIP packets

To extend the NetWare Call filter to also prevent IPX RIP packets from resetting the idle timer or initiating a call, you can define the following additional Output filters:

```
Out filter 07...Generic...Forward=No
Out filter 07...Generic...Offset=0
Out filter 07...Generic...Length=6
Out filter 07...Generic...Mask=ffffffffffff0000
Out filter 07...Generic...Value=ffffffffffff0000
Out filter 07...Generic...Compare=Equals
Out filter 07...Generic...More=Yes

Out filter 08...Generic...Forward=No
Out filter 08...Generic...Offset=24
Out filter 08...Generic...Length=8
Out filter 08...Generic...Mask=ffffffffffffffff
Out filter 08...Generic...Value=ffffffffffff0453
```

Defining Filters and Firewalls

Working with predefined call filters

```
Out filter 08...Generic...Compare=Equals
Out filter 08...Generic...More=No

Out filter 09...Generic...Forward=No
Out filter 09...Generic...Offset=0
Out filter 09...Generic...Length=6
Out filter 09...Generic...Mask=ffffffffffff0000
Out filter 09...Generic...Value=ffffffffffff0000
Out filter 09...Generic...Compare=Equals
Out filter 09...Generic...More=Yes

Out filter 10...Generic...Forward=No
Out filter 10...Generic...Offset=27
Out filter 10...Generic...Length=8
Out filter 10...Generic...Mask=ffffffffffffffff
Out filter 10...Generic...Value=ffffffffffff0453
Out filter 10...Generic...Compare=Equals
Out filter 10...Generic...More=No

Out filter 11...Generic...Forward=Yes
Out filter 11...Generic...Offset=0
Out filter 11...Generic...Length=0
Out filter 11...Generic...Mask=0000000000000000
Out filter 11...Generic...Value=0000000000000000
Out filter 11...Generic...Compare=Equals
Out filter 11...Generic...More=No
```

Defining a SNEP data filter for Ethernet

NetWare's copy-protection scheme makes use of Serialization Number Exchange Protocol (SNEP) packets, which are sent and received by all servers on the network. SNEP packets occur as request/response pairs between servers. When NetWare servers are supported on both sides of the WAN, these packet exchanges can keep an IPX connection active unnecessarily.

This example SNEP filter is intended to be applied as a data filter on the Ethernet interface. To create a SNEP data filter for the Ethernet interface of the Pipeline, create a new Filter profile and define the following Input filters:

```
In filter 01...Generic...Forward=No
In filter 01...Generic...Offset=30
In filter 01...Generic...Length=2
In filter 01...Generic...Mask=ffff000000000000
In filter 01...Generic...Value=0457000000000000
In filter 01...Generic...Compare=Equals
In filter 01...Generic...More=No

In filter 02...Generic...Forward=No
In filter 02...Generic...Offset=33
In filter 02...Generic...Length=2
In filter 02...Generic...Mask=ffff000000000000
In filter 02...Generic...Value=0457000000000000
In filter 02...Generic...Compare=Equals
In filter 02...Generic...More=No

In filter 03...Generic...Forward=Yes
In filter 03...Generic...Offset=0
In filter 03...Generic...Length=0
In filter 03...Generic...Mask=0000000000000000
In filter 03...Generic...Value=0000000000000000
In filter 03...Generic...Compare=Equals
In filter 03...Generic...More=No
```

If you have enough Output filters available in the NetWare Call filter (for example, when you don't extend the filter to include RIP as described in "Extending the predefined filter for RIP packets" on page 6-23, or if you're using NetWare 4.0 or higher and you don't need the predefined SAP filters) you could choose instead to include these SNEP filters as Output filters in the Call Filter.

IP Call filter

The predefined IP Call filter prevents inbound packets from resetting the idle timer. It does not prevent any type of outbound packets from resetting the timer or placing a call.

The IP Call filter contains one Input filter, which defines all inbound packets, and one Output filter, which defines all outbound packets destined for the remote network specified in a Connection or Answer profile in which the filter is applied.

```
In filter 01...Generic...Forward=No
In filter 01...Generic...Offset=0
In filter 01...Generic...Length=0
In filter 01...Generic...Mask=000000000000000000
In filter 01...Generic...Value=0000000000000000
In filter 01...Generic...Compare=Equals
In filter 01...Generic...More=No

Out filter 01...Generic...Forward=Yes
Out filter 01...Generic...Offset=0
Out filter 01...Generic...Length=0
Out filter 01...Generic...Mask=000000000000000000
Out filter 01...Generic...Value=0000000000000000
Out filter 01...Generic...Compare=Equals
Out filter 01...Generic...More=No
```

AppleTalk Call filter

The AppleTalk Call filter instructs the Pipeline to place a call and reset the idle timer based on AppleTalk activity on the LAN, but to prevent inbound packets or AppleTalk Echo (AEP) packets from resetting the timer or initiating a call. The Call filter includes one Input filter and five Output filters.

The Input filter prevents inbound packets from resetting the idle timer or initiating a call. The first two Output filters identify the AppleTalk Phase II AEP protocol, and the next two Output filters identify AppleTalk Phase I AEP protocol. Because More is set to Yes in the first and No in the second filter of

these two pairs, a packet has to meet the criteria defined in both filters to be considered a match. The last Output filter tells the Pipeline to allow all other outbound packets to reset the idle timer or initiate a call.

```
In filter 01...Generic...Forward=No
In filter 01...Generic...Offset=0
In filter 01...Generic...Length=0
In filter 01...Generic...Mask=000000000000000000
In filter 01...Generic...Value=0000000000000000
In filter 01...Generic...Compare=Equals
In filter 01...Generic...More=No
```

```
Out filter 01...Generic...Forward=No
Out filter 01...Generic...Offset=14
Out filter 01...Generic...Length=8
Out filter 01...Generic...Mask=ffffff000000ffff
Out filter 01...Generic...Value=aaaa03000000809b
Out filter 01...Generic...Compare=Equals
Out filter 01...Generic...More=Yes
```

```
Out filter 02...Generic...Forward=No
Out filter 02...Generic...Offset=32
Out filter 02...Generic...Length=3
Out filter 02...Generic...Mask=ffffff0000000000
Out filter 02...Generic...Value=0404040000000000
Out filter 02...Generic...Compare=Equals
Out filter 02...Generic...More=No
```

```
Out filter 03...Generic...Forward=No
Out filter 03...Generic...Offset=12
Out filter 03...Generic...Length=2
Out filter 03...Generic...Mask=ffff000000000000
Out filter 03...Generic...Value=809b000000000000
Out filter 03...Generic...Compare=Equals
Out filter 03...Generic...More=Yes
```

```
Out filter 04...Generic...Forward=No
Out filter 04...Generic...Offset=24
```

```
Out filter 04...Generic...Length=3
Out filter 04...Generic...Mask=ffffff0000000000
Out filter 04...Generic...Value=0404040000000000
Out filter 04...Generic...Compare=Equals
Out filter 04...Generic...More=No

Out filter 05...Generic...Forward=yes
Out filter 05...Generic...Offset=0
Out filter 05...Generic...Length=0
Out filter 05...Generic...Mask=0000000000000000
Out filter 05...Generic...Value=0000000000000000
Out filter 05...Generic...Compare=Equals
Out filter 05...Generic...More=No
```

Display unwanted dial-out packets

A diagnostic option captures and displays packets that cause the Pipeline to dial out. You can then use the information to write data or call filters to prevent the packets from bringing up unwanted connections.

When packets are not captured

If a dial out is initiated for any of the following reasons, the wdDialout option does *not* capture a packet:

- Dial out caused by the Ctrl-D user command
- Dial out caused by callback security
- Dial out on nailed channels
- Dial out caused by NAT (Network Access Translation) acquiring an IP address
- Dial out initiated for IP over X.25, when the X.25 internet profile changes to active and there is data waiting for X.25 to bring up the connection
- Dial out caused by IGMP (Internet Group Management Protocol) multicast forwarding
- Dial out to acquire a DNS address during PPP negotiations

- Dial out in response to a DHCP Discover message
- Dial out caused by the Pipeline sending a DHCP packet for DHCP client processing
- Dial out caused in response to an APP (Ascend Password Protocol) Connect Request message

Turning on the diagnostic option

- 1 Use the DO command D-Diagnostic to open the Diagnostic monitor.
- 2 At the prompt (>) type:

```
help ascend
```

you should see the wdDialout option listed. By default, the option is off.
- 3 To turn the option on, type:

```
wdDialout
```

```
WANDATA dialout display is ON
```

This is a toggle command. Typing it again turns the option off. See the next section for details on how packets are displayed in the diagnostic monitor.
- 4 To exit the diagnostic mode and return to the VT100 interface, type:

```
quit
```

Displaying packets

You can view wdDialout output in the diagnostic monitor. This section shows several examples.

Example 1

In the following example, the Pipeline unit's time and date have not been explicitly set, so the date and time in the captured packet is invalid. The phone number dialed on receipt of this packet is 92233002.

Defining Filters and Firewalls

Display unwanted dial-out packets

```
Date: 01/01/1990.           Time: 00:00:53
Cause an attempt to place call to 92233002
WD_DIALOUT_DISP: chunk 260126 type OLD-STYLE-PADDED.
: 42 octets @ 2C6950
[0000]: ff ff ff ff ff ff 00 c0 7b 61 44 fe 08 06 00 01
[0010]: 08 00 06 04 00 01 00 c0 7b 61 44 fe cc b2 d7 7b
[0020]: 00 00 00 00 00 00 cc b2 d7 13

[0000]: ff ff ff ff ff ff 00 80 c7 5b e9 5b 08 06 00 01
[0010]: 08 00 06 04 00 01 00 80 c7 5b e9 5b cc b2 d7 13
[0020]: 00 00 00 00 00 00 cc b2 d7 16 00 00 00 00 00
[0030]: 00 00 00 00 00 00 00 00 00 00 00 00
```

The type OLD-STYLE-PADDED means that the packet has a 14-byte MAC (Ethernet) header + datagram (ARP request message). The packet contents provide the following information:

destination MAC address	ff:ff:ff:ff:ff:ff	
source MAC address	00:c0:7b:61:44:fe	/* 123 */
arp packet type	08:06	
arp_hrd	00:01	/* Ethernet1 */
arp_prot	08:00	/* IP=0x800 */
arp_hlen	06	/* hlen = 6 */
arp_plen	04	/* plen = 4 */
arp_op	00:01	/* arp ARP_REQ */
arp_sha	00:c0:7b:61:44:fe	/* 123 */
arp_spa	cc:b2:d7:7b	/* 123 */
arp_tha	00:00:00:00:00:00	
arp_tpa	cc:b2:d7:13	/* 19 */

Example 2

In this example, the phone number dialed on receipt of this packet is 92233002. The type OLD-STYLE-PADDED means that the packet has a 14-byte MAC (Ethernet) header + datagram. This is a broadcast IP RWHO message.

```
Date: 01/01/1990.  Time: 00:00:56
Cause an attempt to place call to 92233002
WD_DIALOUT_DISP: chunk 260126 type OLD-STYLE-PADDED.
: 198 octets @ 296810
[0000]: ff ff ff ff ff ff 00 80 c7 5b e9 5b 08 00 45 00
[0010]: 00 b8 0d c3 00 00 3f 11 24 fa cc b2 d7 13 cc b2
[0020]: d7 ff 02 01 02 01 00 a4 e5 8a 01 01 00 00 32 46
[0030]: 5e 26 00 00 00 00 63 6d 61 72 69 6e 65 72 00 00
[0040]: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
[0050]: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
[0060]: 00 00 32 46 4a e3 74 74 79 63 32 00 00 00 72 79
[0070]: 75 00 00 00 00 00 32 46 4b 35 00 00 02 59 74 74
[0080]: 79 63 33 00 00 00 72 79 75 00 00 00 00 00 32 46
[0090]: 4b 39 00 00 00 3d 74 74 79 63 34 00 00 00 72 79
[00a0]: 75 00 00 00 00 00 32 46 4b 3e 00 00 00 97 74 74
[00b0]: 79 70 30 00 00 00 72 79 75 00 00 00 00 00 32 46
[00c0]: 5e 00 00 00 00 01
```

The packet contents provide the following information:

destination MAC address	ff:ff:ff:ff:ff:ff	
source MAC address	00:80:c7:5b:e9:5b	
source IP address	cc:b2:d7:13	/* 204.178.215.19 */
destination IP address	cc:b2:d7:ff	/* 204.178.215.255 subnet broadcast */

Example 3

In this example, the phone number dialed on receipt of this packet is 92233002. The type OLD-STYLE-PADDED means that the packet has a 14-byte MAC header + datagram. This is a unicast IP ICMP echo packet message.

Defining Filters and Firewalls

Display unwanted dial-out packets

```
Date: 01/01/1990.      Time: 00:01:13
Cause an attempt to place call to 92233002
WD_DIALOUT_DISP: chunk 260126 type OLD-STYLE-PADDED.
: 98 octets @ 291EC8
[0000]: 08 00 20 1f 5b ce 00 80 c7 5b e9 5b 08 00 45 00
[0010]: 00 54 0e 09 00 00 ff 01 66 10 cc b2 d7 13 cc b2
[0020]: d7 16 08 00 f5 1b bb 07 98 00 37 5e 46 32 3a 48
[0030]: 0d 00 08 09 0a 0b 0c 0d 0e 0f 10 11 12 13 14 15
[0040]: 16 17 18 19 1a 1b 1c 1d 1e 1f 20 21 22 23 24 25
[0050]: 26 27 28 29 2a 2b 2c 2d 2e 2f 30 31 32 33 34 35
[0060]: 36 37
```

The packet contents provide the following information:

destination MAC address	08:00:20:1f:5b:ce
source MAC address	00:80:c7:5b:e9:5b
source IP address	cc:b2:d7:13 /* 204.178.215.19 */
destination IP address	cc:b2:d7:ff /* 204.178.215.22 */

Example 4

In this example, the phone number dialed on receipt of this packet is 917007337921. Note that there is no MAC header. This is an IPX packet: a Get Nearest Server Request with service type File Server (0004).

```
Date: 01/01/1990.      Time: 00:01:43
Cause an attempt to place call to 917007337921
WD_DIALOUT_DISP: chunk 261022 type IPX.
: 34 octets @ 2C6AA0
[0000]: ff ff 00 22 00 11 00 00 00 00 ff ff ff ff ff ff
[0010]: 04 52 00 00 00 00 00 a0 24 be d5 84 40 09 00 03
[0020]: 00 04
```

The packet contents provide the following information:

```
chksum          ff:ff
packet len      00:22          /* 34 */
Transport Control 00          /* 0 */
packet type     11            /* 17 NCP Packet */
dest network    00:00:00:00
dest Node       ff:ff:ff:ff:ff:ff
dest Socket     04:52          /* SAP */
source network  00:00:00:00:00
source Node     00:a0:24:be:d5:84 /*physical addr of Node*/
Source Socket   40:09          /*4000h-7ffffh Dynamic socket*/
Sap operation   00:03          /* Get Nearest Server Request */
Sap Service Type 0:04          /* File Server */
```

Example 5

In this example, the phone number dialed on receipt of this packet is 92233002.
The type OLD-STYLE-PADDED means that the packet has a 14-byte MAC header + datagram.

```
Date: 01/01/1990.  Time: 02:40:35
Cause an attempt to place call to 92233002
WD_DIALOUT_DISP: chunk 260126 type OLD-STYLE-PADDED.
: 60 octets @ 2AE950
[0000]: 00 80 5f 74 93 d5 00 80 c7 2f 32 4c 00 2a ff ff
[0010]: 00 29 00 11 30 6c 6b 00 00 00 00 00 00 01 04 51
[0020]: 82 c1 b6 bf 00 80 c7 2f 32 4c 40 03 22 22 3f 03
[0030]: 01 00 16 00 02 15 01 ff ff ff ff ff
```

The packet contents provide the following information:

```
destination MAC address    00:80:5f:74:93:d5
source MAC address         00:80:c7:2f:32:4c
chksum                     ff:ff
packet len                 00:29             /*41*/
packet type                11                /*17 NCP Packet */
dest network               30:6c:6b:00
dest Node                  00:00:00:00:00:01
dest Socket                04:51            /* NCP Pkt*/
source network             82:c1:b6:bf
source Node                00:80:c7:2f:32:4c /* addr of src Node */
Source Socket              40:03            /*4000h-7fffh Dynamic socket*/
```

Secure Access Firewalls

Determining if Secure Access is present

All software that includes Secure Access includes the Sec Acc field in the Sys Options menu. If the feature has not yet been enabled, the option is marked as Not Inst. If the feature has been enabled, the option will be marked as Installed.

```
00-100 Sys Options
>Switched Installed^
  Frm Rel Installed
  Sec Acc Installed  V
```

Firewall profiles

When Secure Access Firewall software is present, you can see if any firewalls are in place on your Pipeline by doing the following:

- 1 Open Ethernet > Firewalls > *any profile*.

For example:

```
Name=
Version=
Length
```


- Name specifies the name of the firewall and is originally created using the Secure Access Manager (SAM) graphical user interface.
- Each firewall contains a version number to ensure that any firewall that is uploaded to the router will be compatible with the firewall software on the router. Secure Access Manager (SAM) checks the version number before uploading a firewall. In the event that a router with a stored firewall profile receives a code update that make the existing firewall incompatible, a default firewall is enabled, permitting only Telnet access to the Pipeline. You cannot edit this field.
- Length specifies the length of the firewall uploaded to the Pipeline from Secure Access Manager (SAM), and cannot be edited.

Assigning firewalls to a Connection profile

You can assign firewalls to a Connection profile to filter incoming or outgoing traffic on a WAN connection. Filters assigned to a Connection profile are activated whenever the WAN session comes online.

To assign a firewall to a Connection profile:

- 1 Create a firewall filter using SAM.
- 2 Download it to the Pipeline.
- 3 Open Ethernet > Connections > *any profile* > Session Options.
- 4 Enter the number of the firewall filter you want to use in the Data filter field.
This number is derived from the number in the Firewall menu by adding 100 to the last 2 digits of the firewall index. For example, if the firewall is number 20-503, enter number 103 in the Data Filter field.
- 5 Exit and save the Connection profile.

Assigning firewalls to the Mod Config profile

Firewalls assigned to the Ethernet > Mod Config profile are used to filter incoming or outgoing traffic on the Ethernet interface. Filters assigned to the Mod Config profile are activated as soon as you save the changes to the Mod Config profile.

To assign a firewall to the Mod Config profile, do the following:

- 1 Create a firewall filter using SAM.
- 2 Download it to the Pipeline.
- 3 Open Ethernet > Mod Config > Ether Options.
- 4 Enter the number of the firewall filter you want to use in the Filter field.
This number is derived from the number in the Firewall menu. For example, if the firewall is number 20-503, enter number 103 in the Data Filter field.
- 5 Exit and save the profile.

Filter persistence

A Filter persistence parameter is present in Connection profiles of all Pipelines that support Filter Profiles. The Filter Persistence parameter must be set to Yes to allow a connection's firewalls to persist when the connection is torn down, such as by connection timeout. The default is No, implying that, by default, connection firewalls do not persist when a call is terminated.

Note: Typically a firewall will persist for about an hour after its associated connection has been torn down.

Background on firewall and filter persistence

The idea of filter persistence is intended to allow a Pipeline to preserve its filter or firewall specifications throughout the lifetime of its connections.

Firewalls differ from filters in that firewalls are designed to alter their behavior as traffic passes through them, but filters remain unchanged through their lifetimes.

Filters provide for the construction and destruction of filters whenever the state of a connection changes, which causes the Pipeline to create and destroy filters during connection state changes without any reference to the state of the filters.

When Secure Access Firewalls are present, it is necessary to preserve the firewall state across the many transitions that connections may experience. Where filters can be built or destroyed at any time to accommodate changes due to Multilink and idle-inactivity conditions, firewalls cannot.

A persistent filter or firewall is maintained even when its associated connection becomes inactive. Additionally, the filter or firewall can be applied when an additional session becomes associated with a connection, as is the case with additional channels of an MPP connection.

Note: Firewalls need to use persistence to work correctly, but filters do not need to use persistence to work as designed.

Filter persistence and Connection profiles

Connection profiles describe different contact sites. Perhaps, for a small office, one profile would apply to a corporate home office, and another profile would apply to an Internet service provider. In each case, the Pipeline user would like to use the Secure Access Firewall capability to prevent unauthorized incursions into the local network by others.

With dial-on-demand and automatic call timeout, the dynamic firewall capabilities of Secure Access Firewall would prevent in-progress TCP sessions (such as Telnet or Rlogin) from proceeding after a call termination and restart (due to inactivity, for example). Without persistence, a new firewall is constructed when a call starts up with no knowledge of any TCP sessions in progress, and consequently would block packets for those sessions when starting the line back up. This has the effect of rendering the in-progress Telnet (or Rlogin, etc.) sessions inoperative, possibly destroying work in progress that is dependent on them.

Filter persistence is a way to tell the Pipeline to keep a firewall around even after the call is terminated. When a new call is placed to (or is received from) the same station, the Pipeline remembers the original firewall and uses it as if the call had never been terminated. Thus, the user can continue working without loss.

Conversely, there may be times when a single Connection profile is used for several different sites. This might be the case if you use the same Connection Profile to describe multiple different callers. In this case, you do not want the filters and firewalls to be persistent, since the Pipeline cannot know if calls are arriving from the same users.

Setting Up Pipeline Security

This chapter includes the following topics:

Recommended security measures.	7-1
Pipeline Security profiles	7-7
Connection security	7-11
Using filters to secure the network.	7-16
Using security cards	7-17

Recommended security measures

When the Pipeline is shipped from the factory, its security features are all set to defaults that enable you to configure and set up the Pipeline without any restrictions. Before you make the Pipeline generally accessible, you should change these default security settings to protect the configured unit from unauthorized access.

You should set these important security features before putting the Pipeline online:

- Change the Full Access security level password.
A user who knows the password to the Full Access level will be able to perform any operation on the Pipeline, including changing the configuration. The Full Access password is set to “Ascend” by default, and you should assign your own password. (For instructions, see “Changing the Full Access security level password” on page 7-3.)

Setting Up Pipeline Security

Recommended security measures

- Activate the Full Access security level.

After you change the password, activate the Full Access security level for your own use in performing the rest of these basic security measures. (For instructions, see “Activating the Full Access security level” on page 7-4.)

- Make the default security level very restrictive.

The Pipeline provides terminal services via Telnet. Any user who Telnets to the unit is assigned the default security level, which is initially without restrictions. You should turn off all privileges in the Default security profile. (For instructions, see “Making the Default security level restrictive” on page 7-4.)

- Assign a Telnet password.

Until you assign a Telnet password, any local user who has the Pipeline unit’s IP address can Telnet into it. Once you assign the password, all incoming Telnet sessions (from the local network or across the WAN) will be prompted to enter that password. (For instructions, see “Assigning a Telnet password” on page 7-5.)

- Change the SNMP community strings.

The Pipeline supports SNMP traps, which allows it to send alarms, report on call details, and send other management information to an SNMP management station without being polled. The Pipeline default read and write community strings should be changed to prevent unauthorized access to the Pipeline by an SNMP management station. (For instructions, see “Changing the SNMP read and write community string” on page 7-5.)

- Require profiles for incoming connections.

The Pipeline unit’s Answer profile can be used to build unrestricted connections (connections for which no name or password is required). Although some sites allow this type of connection, many do not. You should restrict incoming connections to those with a configured profile. (For instructions, see “Requiring profiles for incoming connections” on page 7-7.)

- Turn off ICMP Redirects.

To secure the Pipeline unit’s IP routes, you should configure the unit to ignore ICMP (Internet Control Message Protocol) Redirect packets. (For instructions, see “Turning off ICMP redirects” on page 7-7.)

Changing the Full Access security level password

The Full Access security profile is intended to provide unrestricted access to the Pipeline. This is the “super-user” profile that enables you to configure, dial-up remote locations, reset the unit, upgrade system software, and so forth.

Note: Write down and save the Full Access password in a safe place. Make sure when you open the Full Access profile that you do not turn off the Edit Security privilege, or you will be unable to edit privileges when Full Access is activated.

To change the Full Access password:

- 1 Open the System > Security menu.

```
00-300 Security
>00-301 Default
00-302
00-303 Full Access
```

- 2 Open the Full Access profile.

```
00-303 Full Access
Name=Full Access
>Passwd=ascend
Operations=Yes
Edit Security=Yes
Edit System=Yes
Field Service=Yes
```

- 3 Open the Passwd parameter and specify a new password, then press Enter.

For example:

```
Passwd=my-password
```

Note: Passwords are case-insensitive. A user can specify the password “my-password” as “My-Password” or “MY-PASSWORD” and the Pipeline accepts it.

- 4 Leave all other privileges enabled.

Note: Do not turn off the Edit Security privilege in this profile!

- 5 Close the Full Access profile.

Now only users who have the password you assigned will be able to activate the Full Access security level.

Activating the Full Access security level

To activate the Full Access profile, do the following:

- 1 From the VT100 menus, press Ctrl-D to open the DO menu, and then press P (or select P=Password).

```
DO...
>0=ESC
P=Password
```

- 2 In the list of security profiles, select Full Access. The Pipeline prompts for the password.

```
00-300 Security
Enter Password:
[ ]
```

Type the password you specified in the Full Access profile and press Enter.

A message states that the password was accepted and the Pipeline is using the new security level. If the password you enter is incorrect, you are prompted again to enter the password.

Making the Default security level restrictive

The Default security level is always assigned to all users who Telnet into the unit or access the terminal server interface in another way, and it is activated for the console whenever the unit is reset. You cannot change the name of the Default security profile or assign a password to it, but you should turn off its operations privileges.

To set the default security level to allow read-only privileges:

- 1 Open the System > Security > Default profile.
- 2 Restrict the Operations privilege.

For example:

Operations=No

When you restrict this privilege, all other privileges are N/A.

3 Close the Default profile.

Once set, users who access the Pipeline terminal server will be unable to make any changes to its configuration or perform restricted operations. For all users with the default security level, passwords (including the null password) will be hidden by the string *SECURE* in the Pipeline user interface.



Caution: Resetting or powering the unit on and off activates the new, restrictive Default profile. You will not be able to perform any configuration tasks until you activate and supply the password for the Full Access profile. Use the default password “Ascend” to access the Full Access profile.

Assigning a Telnet password

Assign a Telnet password to prevent unauthorized Telnet sessions. The Telnet password can be up to 20 characters in length.

To assign a Telnet password:

- 1** Open the Ethernet > Mod Config > Ether Options.
- 2** Enter a Telnet password up to 20 characters long.

For example:

```
Telnet PW=telnet-pwd
```

- 3** Close the Ethernet profile.

Now any user who opens a Telnet session to the Pipeline will be prompted to supply this password.

Changing the SNMP read and write community string

SNMP community strings are identifiers that SNMP-manager applications must specify before they can access the Management Information Base (MIB). The Pipeline has two community strings:

- Read Comm

Setting Up Pipeline Security

Recommended security measures

The read community string enables an SNMP manager to perform read commands (for example, Get and Get next) to request specific information.

- R/W Comm

The read-write community string enables an SNMP manager to perform both read and write commands (for example, Get, Get next, and Set), which means the SNMP application can access management information, set alarm thresholds, and change some settings on the Pipeline.

To enable SNMP set commands, enter a Read Community password that must be known by the SNMP manager in order to read the Pipeline settings, and enter a password for Read/Write Community that must be known by the SNMP manager in order to change the settings.

- 1 Open the Ethernet > Mod Config > SNMP Options submenu.

```
Read Comm=public
>R/W Comm Enable=Yes
R/W Comm=write
```

The default Read Community name is “public” and the default Read/Write Community password is “write”.

- 2 Enter up to 16 alphanumeric character in the Read Comm parameter.

For example:

```
Read Comm=name
```

- 3 To enable the use of SNMP set commands, set the Read/Write Community string parameter to yes.

For example:

```
R/W Comm Enable=Yes
```

When the value is No, the R/W Comm parameter is N/A.

- 4 Enter up to 16 alphanumeric characters in the R/W Comm parameter.

For example:

```
R/W Comm=unique-string
```

- 5 Close and save the profile.

Note: To use a Set command (or use the Java-based Pipeline Configurator to update the unit), you must know the R/W value, and R/W Comm Enable must be set to Yes.

Requiring profiles for incoming connections

There are many authentication measures you can set for incoming connections. At the most basic level, you can configure the Pipeline to reject all incoming calls that don't have a Connection profile.

To require configured profiles for all incoming connections:

- 1 Open the Ethernet > Answer profile.
- 2 Specify that a matching profile is required for incoming calls.

For example:

```
Profile Req'd=Yes
```

- 3 Close and save the profile.

(For more information about securing incoming connections, see "Connection security" on page 7-11.)

Turning off ICMP redirects

Internet Control Message Protocol (ICMP) was designed to dynamically find the most efficient IP route to a destination. ICMP Redirect packets are one of the oldest route discovery methods on the Internet and one of the least secure. It is possible to create counterfeit ICMP Redirects and change the way a device routes packets. If the Pipeline is routing IP, you should turn off ICMP redirects.

To configure the Pipeline to ignore ICMP redirect packets, do the following:

- 1 Open the Ethernet > Mod Config profile.
- 2 Turn off ICMP redirects.

For example:

```
ICMP Redirects=Ignore
```

- 3 Close and save the profile.

Pipeline Security profiles

When the Pipeline is shipped from the factory, its security privileges are open to enable you to configure and set it up without any restrictions. (For recommended

settings for the two predefined Security profiles, see “Recommended security measures” on page 7-1.)

Default security level

The Pipeline has three possible security levels, including the default. The Default security profile has no password. This security level is always activated for all users who Telnet into the unit or access the terminal server interface in another way. The Default security level is activated for the console whenever the unit is reset, so that the privileges enabled in the Default profile are generally available. Set System > Security > Default profile, Operations=No to prevent unauthorized changes to other settings.

Security profile passwords

Passwords are case-insensitive in the Pipeline. If you specify the password “my password,” the Pipeline accepts that string in any case combination (such as “My-Password” or “MY-PASSWORD”).

Users who do not have Edit Security privileges, described next, can see the Pipeline menus, but all passwords are displayed as *SECURE* instead of the actual password. If a user has Edit Security privileges, passwords in Security profiles can be seen and changed.

Security privileges

In addition to Default security, there is an additional Security profiles you can customize to include any combination of the following privileges:

- Operations
If Operations=Yes, users can change parameter settings and access most DO commands, which are manual commands used to change security levels or manually dial or clear calls. (To learn more about DO commands, see the *Reference Guide* chapter on using the Do commands.)
- Edit Security
If Edit Security=Yes, users can edit Security profiles. All passwords in Security profiles are visible as text. This is the most powerful privilege you can assign, because it allows users to change their own privileges at will.

When Edit Security=No, all passwords are hidden by the string
“*SECURE*.”

- Edit System
If Edit System=Yes, users can edit the System profile and other system-wide settings.
- Field Service
If Field Service=Yes, users can perform field service operations, such as uploading new system software to the Pipeline unit. Field service operations are special diagnostic routines not available through Pipeline menus.

For complete information on each parameter, see the *Reference Guide*.

Using the Full Access profile

The Full Access profile should be reserved for the super-user login: yourself and anyone else who will be reconfiguring the Pipeline, testing lines, dialing remote locations, resetting the unit, and upgrading system software.

Note: Be sure you write down the new Full Access password and store it in a safe place. If you restrict all other levels and then forget the Full Access password, you will need to call Customer Support to access the unit.

The default settings for the Full Access profile are as follows:

Name=Full Access
Passwd=Ascend

Note: You should change this default password, as described in
“Recommended security measures” on page 7-1

Operations=Yes
Edit Security=Yes

Note: Do not turn off the Edit Security privilege, or you will be unable to edit privileges when Full Access is activated!

Edit System=Yes
Field Service=Yes

When you log into the Pipeline, you will only be able to view settings, because the Default profile will be active. To make any changes or perform any administrative tasks, you need to activate the Full Access profile in the DO menu.

Setting Up Pipeline Security

Pipeline Security profiles

(To learn more about DO commands, see the *Reference Guide* chapter on using the Do commands.)

- 1 Press Ctrl-D to open the DO menu, and then press P (or select P=Password).

```
DO...
>0=ESC
P=Password
```

- 2 Open System > Security > Full Access.
The Pipeline prompts for the password.
- 3 Type the password for the Full Access profile and press Enter.

Defining a second Security profile

If you do not want other users to change the Pipeline configuration profiles or perform administrative tasks in the Pipeline, you do not need to define any Security profiles beyond Default and Full Access. However, you can define an additional security profile, as described below.

To define a Security profile:

- 1 Open the System > Security > *unnamed profile*.
- 2 Specify a name for the profile (up to 16 characters).

For example:

```
Name=Calabasas
```

- 3 Specify a new password, and then press Enter.

```
Passwd=*SECURE*
```

As soon as you press Enter, the Pipeline hides the password string you specified by displaying the string *SECURE*.

- 4 Set the privileges for this profile.

For example:

```
Name=Calabasas
Passwd=*SECURE*
Operations=Yes
Edit Security=No
Edit System=No
Field Service=No
```

- 5 Close and save the profile.

Connection security

Connection security has two levels: caller authentication regulating authorized access, and network security preventing unauthorized wide-area network access.

All authentication relies on the Pipeline finding a matching profile to verify information presented by the caller.

- Authentication mechanisms
 - Password authentication, such as PAP, CHAP, or MS-CHAP, requires a name and password from the caller. Additionally, CHAP encrypts the password data.
 - Calling-line ID (CLID) authentication verifies that the call is coming from the expected phone number.
 - Called number (Called #) is similar to authentication by CLID, but it authenticates on the number called into rather than the number originating the call. In the Connection profile, the Called # parameter is almost identical to the Dial # parameter, but uses a number without a trunk group or dialing prefix prepended.
 - Callback authentication instructs the Pipeline to hang up and call back before performing password authentication. Callback provides the highest level of control, assuring that incoming calls are coming from a known user or network.

Note: Any form of authentication requires a configured profile. See “Requiring profiles for incoming connections” on page 7-7 for details on configuring the Pipeline to always require a matching profile, regardless of whether authentication is enforced.

- Network security
 - Filters are one of the most effective methods of protecting your site from unwanted WAN access. Filters are described briefly in this chapter; see Chapter 6, “Defining Filters and Firewalls,” for full details.

Authentication protocols

Password Authentication Protocol (PAP) and Challenge Handshake Authentication Protocol (CHAP) require Point-to-Point Protocol (PPP) encapsulation. These authentication protocols apply to PPP, Multilink PPP (MP), and Multichannel PPP (MP+) connections to the Pipeline. Both sides of the connection must support the same protocol.

PAP provides a simple way for a peer to establish its identity in a two-way handshake when initially establishing a link. It sends passwords in the clear, so it is not a very strong authentication method. PAP provides baseline security when your system interoperates with equipment from other vendors.

CHAP is a stronger authentication method than PAP. During the establishing of the initial link, CHAP verifies the identity of a peer through a three-way handshake. It sends passwords encrypted by means of a one-way hash function. This use of an incrementally changing identifier and a variable challenge value protects against playback attack.

MS-CHAP is supported to work with DES and MD4 encryption in Windows NT environments only. The Pipeline can authenticate a Windows NT system and a Windows NT system can authenticate a Pipeline.

Note: In addition to this type of authentication, there are other parameters, such as Telco and Session options, that affect whether the Pipeline is able to build the connection. For example, if the AnsOrig parameter is set to prevent incoming calls, the Pipeline will never reach the stage of authenticating an incoming call using that profile.

Name and password verification

During authentication, the calling device often requires the Pipeline unit's name and password as well. The Pipeline name is specified in the System profile. The Send PW parameter is a password sent to the calling device.

If the Ethernet > Answer > Recv Auth parameter is set to Either, the Pipeline uses PAP, CHAP, or MS-CHAP, depending on what the caller supports. If it is set just to a specific authentication protocol, the Pipeline rejects any password not sent with the assigned authentication protocol.

When the Pipeline receives a PPP call, it tries to match the caller's name and password to a Connection profile. If the Pipeline doesn't find a matching profile, it ends the call. If the Pipeline finds a matching profile, it authenticates the call and establishes the connection.

When an IP routing connection is being authenticated, the IP address is verified as part of the PPP negotiation before a call is established.

The Ethernet > Answer > PPP options > Route IP parameter must be set to Yes.

If the caller's PPP software presents an IP address, the Pipeline must find a Connection profile that matches that address using Ethernet > Connections > *any profile* > IP options > LAN Adrs parameter, which must contain a matching IP address. Otherwise, it ends the call without completing PAP or CHAP authentication. If it finds a profile, it authenticates the connection, and then establishes the connection.

Calling-line ID authentication

Calling-line ID (CLID) requires the phone number of the calling device. CLID authentication ensures that the incoming call originates from a known phone number. Id Auth in Ethernet > Answer must be set to Yes, and you must set a value in the Calling # parameter of the matching Connection profile.

When CLID authentication is required, if the calling number is not recognized, the Pipeline hangs up. CLID authentication occurs first, before any name or password comparison.

Note: In some installations, the WAN provider might not be able to deliver CLID information, or individual callers might choose to block Caller ID. In addition, CLID is not available without end-to-end ISDN service on the call and Automatic Number Identification (ANI) from your WAN provider. Ask your WAN provider whether the calling-party number is conveyed by the network to the receiving party. In some cases, the network does not deliver the calling-party number, such as when the Pipeline is behind one or more PBXs.

Setting Up Pipeline Security

Connection security

The Id Auth parameter in the Answer profile can be set to the following values:

Ignore	Ignore indicates that calling-party information is not required for authentication.
Prefer or Called Prefer	Prefer specifies that whenever CLID is available, the calling-party's phone number must match the Calling # parameter before answering the call. If CLID information is not available or if the Pipeline cannot find a match to a calling number, the Pipeline applies authentication using the Recv Auth or Password Req parameters. Called Prefer is the same as Prefer except that the called number, rather than the calling number is preferred.
Require or Called Require	Required indicates that the calling party's phone number must match the value of the Calling # parameter before the Pipeline can answer the call. If CLID information is not available, the Pipeline does not answer the call. Called Require is the same as Require except the called number, rather than the calling number is required.

Note: Fallback is listed, but is not currently available for the Pipeline.

Settable disconnect cause codes for CLID authentication

When Caller ID authentication fails in an ISDN connection, the Pipeline sends a Disconnect message. The Cause Element in the Disconnect message can give an idea of why the CLID authentication failed. You can set the Disconnect cause code for CLID authentication failures to "User Busy" or "Normal call clearing."

To set the Disconnect Cause value, open the Ethernet > Mod Config > Auth profile. For example:

```
Auth...  
CLID Fail Busy=No  
APP Server=No  
APP Host=N/A  
APP Port=N/A
```

Set the CLID Fail Busy parameter to Yes to make the disconnect message “User Busy;” set it to No for the message “Normal call clearing,” which is the default.

Callback security

Callback security instructs the Pipeline to hang up on an incoming caller and then immediately initiate a call to that destination.

To use Callback security, set the following parameters:

- Ethernet > Connections > *profile* > Calling # and Dial #
Callback ensures that the connection is made with the number specified in the Calling # parameter.
- Ethernet > Connections > *profile* > Telco Options > Callback=Yes
- Ethernet > Connections > *profile* > Telco Options > AnsOrig=Both
When setting Callback=Yes, you must also set AnsOrig=Both, because the Connection profile must answer the call *and* call back the device requesting access. Similarly, the calling device must be able to dial out to *and* accept incoming calls from the Pipeline.

Note: For units whose Call Type=Nailed, indicating a leased line, Callback is not supported.

To set callback security:

- 1 Open Ethernet > Connection profile.
- 2 Specify the number the Pipeline needs to dial to reach the remote device.
For example:
`Dial #=555-1213`
- 3 Specify the number the remote device uses to return the call to the Pipeline.
For example:
`Calling #=555-1214`
- 4 Open the Telco Options submenu.
- 5 Turn on callback security.
For example:
`Callback=Yes`
`AnsOrig=Both`

- 6 Close and save the profile.

Expect callback support

If Ping or Telnet attempt to reach a far end that is using callback security, it causes a problem. Ping and telnet try continuously to open a connection and reject the return callback because the process is already trying to establish a connection.

To remedy the situation, set Expect Callback to Yes. This puts the number of *any* far end that does not connect (for any reason) on a list that disallows calls to that destination for 90 seconds. This gives the far end an opportunity to complete the callback.

Note: Expect Callback should only be set to Yes in dialout profiles.

Set Expect Callback to Yes by doing the following:

- 1 Open Ethernet > Connections > *any profile* > Telco Options.
- 2 Set Exp Callback to Yes.

Now if an outgoing connect from that profile fails (for any reason), you will be forced to wait 90 seconds before attempting to connect again.

Using filters to secure the network

Network security is related to packets coming in from any wide-area network (WAN) connection. The most direct method of securing the network is with filters.

Note: For recommendations about ICMP Redirect packets, see “Recommended security measures” on page 7-1.

Network security filters are data filters, which may be applied to incoming or outgoing data streams, or both. Data filters can prevent certain packets from reaching the local network or going out from the local network to the WAN. For example, you can use data filters to drop packets addressed to particular hosts, or prevent certain types of packets from reaching the local network.

Filters can also be used to prevent remote users from accessing information on your local network, even if they know how to “spoof” a local source address that would enable them to get past a filter. For example, you can define a filter that drops inbound packets whose source address is on the local network or the loopback address.

Each filter consists of an ordered list of conditions (“rules”) based on either IP-specific or protocol-independent information. For an IP filter, you can filter packets based on any combination of the following elements:

- Source address
- Destination address
- Protocol number
- Source port
- Destination port
- A flag indicating if a TCP session is established

For a protocol-independent filter, you can specify data values and masks that the Pipeline uses when determining whether to drop or forward packets.

(For information about how to organize and create Filter profiles, refer to Chapter 6, “Defining Filters and Firewalls.”)

Using security cards

A secure network site can be set up to change its password after a number of minutes or hours. An external authentication server such as a Security Dynamics (ACE) or Enigma Logic (Safeword) server changes the password and relies on a combination of a Personal ID (PIN) and a code generated by security card that must be in the possession of the user. A liquid crystal display on the security card shows the code that enables access to the secure network only at that time.

For secure sites the Pipeline is a client of a central-site device, such as MAX 4000, which acts as a network access server (NAS). The NAS is a client of a RADIUS server, which in turn is a client of the ACE or Safeword server.

Figure 7-1 shows one example security card environment. The user dialing in through a Pipeline unit is a client of the Pipeline, which in turn is a client of the

Setting Up Pipeline Security

Using security cards

MAX (acting as the NAS). The NAS requests authentication from the RADIUS server, which in turn contacts the external server.




Figure 7-1. RADIUS acting as client of ACE or Safeword server

When a user initiates a login to a secure site, the following events occur:

- 1 The calling unit (for example, a Pipeline) calls a NAS (the MAX).
- 2 The NAS requests authentication of the call from the RADIUS server.
- 3 The RADIUS server forwards the request to an ACE or Safeword server.
- 4 The ACE or Safeword server sends a challenge message (which may confirm a null challenge) back through the RADIUS server and the NAS to the calling unit.
- 5 A user on the remote network responds to the challenge message with the current password, which is displayed on the security card.

If the user enters the correct password, network access is established.

If 60 seconds pass without a response to the challenge message, the call is dropped.

If the user enters an incorrect password, the ACE or Safeword server sends another challenge. After three incorrect passwords, the call is dropped.

Supporting outbound security card calls

The instructions in this section explain how you configure the Pipeline to place calls to a remote NAS and to handle password challenges when they are returned from the NAS.

For the Pipeline to place calls to a NAS at a secure site, it needs the appropriate Connection profile specifying a token-based authentication mode.

The authentication mode configured in the Pipeline affects how the token passwords are transmitted and how the dial-in user is affected by channels being added to an established session.

The Pipeline requests the authentication mode with which it is configured, but the RADIUS daemon and user profile accessed by the answering NAS determine which mode will actually be used.

Requesting PAP-TOKEN mode

PAP-TOKEN is the default authentication mode used when the RADIUS profile has a password of ACE or Safeword. It is an extension of PAP authentication.

When PAP-TOKEN mode is in use, the dynamic password (or code) supplied by the user's security card is sent in the clear (via PAP). This does not cause a serious security risk because the password expires every 60 seconds (or at some other very short interval of time).

The response to the initial password challenge authenticates the base channel of the call. If bandwidth requirements cause another channel to come up, the user is challenged for a password whenever a channel is added to a call.

Parameters used to configure the calling unit for PAP-TOKEN are set in the following menus:

```
Ethernet
Connections
  profile
    Encaps options...
      Send Auth=PAP-TOKEN
      Send PW=*SECURE*
```

The Send Auth parameter specifies the authentication mode requested by the caller (PAP-TOKEN). The Send PW password is sent as part of the initial session negotiation. If the session then presents a password challenge, the user enters the password obtained from the security card.

Requesting PAP-TOKEN-CHAP mode

PAP-TOKEN-CHAP authenticates additional channels using CHAP. If it is specified in the Send Auth parameter, but the RADIUS profile at the far end is not set up for PAP-TOKEN-CHAP, then PAP-TOKEN is used instead.

The dynamic password supplied by a user authenticates the base channel of the call. It is sent in the clear (via PAP). When the Pipeline adds additional channels to the call, PAP-TOKEN-CHAP uses CHAP authentication for the new channels. CHAP sends encrypted passwords, so it can take the auxiliary password from the Aux Send PW parameter and transmit it securely.

The following parameters are used to configure the calling unit:

```
Ethernet
Connections
profile
Encaps options...
Send Auth=PAP-TOKEN-CHAP
Send PW=*SECURE*
Aux Send PW=*SECURE*
```

The Send Auth parameter specifies the authentication mode requested by the calling unit (PAP-TOKEN-CHAP in this case). The Send PW password is sent as part of the initial session negotiation. If the session presents a password challenge, the user enters the password generated by the security card.

The Aux Send PW parameter is sent via CHAP for authenticating additional channels; additional entries derived from the security card are not required.

Requesting CACHE-TOKEN

CACHE-TOKEN uses CHAP and caches the initial password for re-use in authenticating channels as they are added to the call. The RADIUS profile at the far end must be set up with appropriate attributes that specify how long the token will be cached.

The following parameters are used to configure the calling unit:

```
Ethernet
  Connections
    profile
      Encaps options...
      Send Auth=CACHE-TOKEN
      Send PW=*SECURE*
```

The Send Auth parameter specifies the authentication mode requested by the calling unit (CACHE-TOKEN). The Send PW password is sent as part of the initial session negotiation. Then, the user is prompted for a token password to authenticate the base channel of the call via CHAP. If the RADIUS server has been configured correctly, it caches that encrypted password for the specified period, or for the specified amount of idle time during the connection. When channels are added to the call or when a new call is made, it uses the cached password to authenticate the connection.

Configuring the Pipeline to recognize the APP Server utility

The Ascend Password Protocol (APP) Server utility lets users respond to token password challenges received from a remote network access server (NAS) from a PC or UNIX host on the local network. To enable the utility, you need to configure the Pipeline to communicate with the host running APP. (For information about obtaining and setting up the APP Server, see “APP Server installation and setup” on page A-2.)

APP is a User Datagram Protocol (UDP) whose default port is 7001. The communication between the Pipeline and the host running the APP Server may be unicast, where both the Pipeline and the host have an IP address; or broadcast, where the host may not have an IP address.

The parameters used to associate the APP server with the Pipeline are:

```
Ethernet
  Mod Config
    Auth...
      APP Server=Yes
      APP Host=10.65.212.1
      APP Port=7001
```


Setting Up Pipeline Security

Using security cards

To set up the Pipeline to communicate with the APP Server utility, do the following:

- 1 Open the Ethernet > Mod Config > Auth menu.

- 2 Set the APP Server parameter to Yes.

For example:

```
APP Server=Yes
```

This enables the Pipeline to communicate password challenges to the host running the APP Server utility.

- 3 Specify the IP address of the host running the APP Server utility.

For example:

```
APP Host=10.65.212.1
```

If the host obtains its IP address from a BOOTP or DHCP server, or if it has no IP address, specify the IP broadcast address of 255.255.255.255.

- 4 Specify the UDP port to use for communicating with the APP host.

For example:

```
APP Port=7001
```

7001 is the default UDP port for the APP Server. The Pipeline and the host running the APP Server utility must agree on the UDP port number. If you use a port number other than 7001, be sure to specify the UDP port number in the APP Server utility (DOS), the WIN.INI file (Windows), or the /etc/services file (UNIX).

- 5 Close the Ethernet profile.

Invoking password mode in the Pipeline

If required, you can also bring up a connection to a secure site when connecting to a remote site using the DO menu, or by dialing the remote NAS via modem.

To invoke password mode in a terminal server session, do the following:

- 1 At the terminal server prompt, type the following:

```
set password
```

The following message is displayed:

```
Entering Password Mode...
```

The prompt changes to the following:

[^C to exit] Password Mode>

- 2 Dial the remote site using any commands you normally use to make the call.

Note: When connecting via modem, skip this step.

- 3 While the connection is being negotiated, the remote NAS returns a challenge prompt similar to the following:

```
From: hostname
0-Challenge: challenge
Enter next password:
```

- 4 The *hostname* is the name of the NAS you called. (Not all systems respond with their host name.)

Enter the password from your security card at the challenge prompt. If the password is:

- Entered correctly, the connection is established to the secure network.
- Entered incorrectly, the challenge prompt is displayed again up to three times.
- Not entered within 60 seconds, the login attempt times out.

If the Send Auth parameter is configured incorrectly, no challenge prompt appears, or you will get an error message such as the following:

```
From: hostname
Received unexpected PAP Challenge!... check PPP Auth Mode
```

- 5 To return to normal terminal server operations, press **Ctrl-C** at the Password Mode prompt.

Pipeline System Administration

8

This chapter includes the following topics:

Overview of administration functions	8-1
Activating administrative privileges	8-3
Configuring administration options	8-4
Using the Pipeline status windows	8-9
Performing system administration operations	8-10
Using the terminal server interface	8-17
Accessing a local Pipeline via Telnet	8-21

Overview of administration functions

The Pipeline provides the following administrative functions:

Security profiles These use password security to protect the unit from unauthorized access. (See “Activating administrative privileges” on page 8-3 and Chapter 7, “Setting Up Pipeline Security.”)

System admin commands	These include commands for rebooting, saving or restoring configuration information, upgrading system software, and viewing statistics and other conditions and settings. (See “Performing system administration operations” on page 8-10. Also see Appendix C, “Upgrading system software.”)
DO commands	Pressing Ctrl-D in the VT100 interface displays the DO menu, which contains commands for changing security levels in the Pipeline, or manually dialing or clearing a call. When full access (or another appropriate security level) has been activated, you can perform all DO commands as well as other administrative operations. (Also see “DO Command Reference” in the <i>Reference Guide</i> .)
Terminal server	The command-line interface provides commands for testing a connection, checking routing tables and other configuration parameters, or configuring far-end Ascend units across the WAN. Many of these commands are related to system administration. (See “Using the terminal server interface” on page 8-17. Also see “Terminal Server Commands” in the <i>Reference Guide</i> .)
Status windows	The status windows in the VT100 interface provide information about what is currently happening in the Pipeline. For example, status windows display the last 31 system events, statistics about the currently active session, the software version loaded on the unit, the hardware configuration, and other information. (See “Status Window Reference” in the <i>Reference Guide</i> .)
Syslog	If a Windows or UNIX host on the local network is running the Syslog daemon, you can configure the Pipeline to write log messages to an ASCII file on that host. (See “Configuring the Pipeline to interact with syslog” on page 8-5.)

SNMP management	<p>The Pipeline supports SNMP on a TCP/IP network. An SNMP management station that uses the Ascend Enterprise MIB can query the Pipeline, set some parameters, sound alarms when certain conditions appear in the Pipeline, and so forth. An SNMP manager must be running on a host on the local IP network, and the Pipeline must be able to find that host, either via static route or RIP.</p> <p>In addition, SNMP has its own password security, which you should set up to protect the Pipeline from being reconfigured from an SNMP station.</p>
Remote management via Telnet	<p>The Pipeline can be remotely configured and managed by establishing a Telnet session from any Telnet workstation and viewing the configuration menus in a Telnet VT100 window. You can use this feature to manage the Pipeline from a local or remote computer. You can also use it to manage remote Ascend units, such as the Pipeline. From a Telnet session you can perform all of the configuration, diagnostic, management, and other functions that could be performed from a computer connected to the Pipeline Terminal port. See “Using the terminal server interface” on page 8-17.</p>

Activating administrative privileges

This section assumes that you have taken the recommended steps to secure the Pipeline box, as described in Chapter 7, “Setting Up Pipeline Security.”

After you have taken the recommended steps, you cannot perform any system administration operations without first supplying the required password. To specify that password:

- 1 Press Ctrl-D to display the DO menu.

```
DO...
>0=ESC
P=Password
E=Termsrv
```

D=Diagnostics

- 2 Press P (or select P=Password) to invoke Password command.
A menu of Security Profiles opens.
- 3 Select Full Access.
The Pipeline prompts for the password for the Full Access profile.

00-300 Security
Enter Password:
[]

Press > to accept
- 4 Type the password and press Enter to accept it.
If you enter the right password, a message states that the password was accepted and the Pipeline is using the new security level. If the password you enter is incorrect, you are prompted again to enter the password.

Configuring administration options

This section describes the following system administration configurations:

- Setting system values
- Specifying administrative information in the System profile
- Setting the Telnet password
- Configuring the Pipeline to interact with a Syslog daemon

Setting system values

The system name is used in negotiating bridged PPP connections. To set the Pipeline unit's system name:

- 1 Open the System > Sys Config menu.
For example:

Name=LAB10GW
Location=LAB10
Contact=MIS

Term Rate=9600
Console=Standard
Remote Mgmt=No

- 2 Specify a system name up to 16 characters long.
- 3 Enter the physical location of the Pipeline.
You can enter up to 80 characters. An SNMP manager can read this field, but its value does not affect the operation of the Pipeline.
- 4 Specify a person to contact in case of error conditions.
You can enter up to 80 characters. An SNMP manager can read this field, but its value does not affect the operation of the Pipeline.
- 5 Specify the data transfer rate of the Pipeline Terminal port.
The default 9600 is appropriate if you are accessing the VT100 interface from a PC connected to the Pipeline Terminal port. If you are managing a remote Ascend unit, you may want to increase the baud rate on the local terminal to a higher speed for improved performance.

Note: Make sure the Term Rate setting matches the speed configured for your Com Port.
- 6 Specify the type of console interface to be displayed at power-up.
Currently the only supported value is “standard.”
- 7 Specify whether a remote device (across the WAN) will be allowed to operate the Pipeline.
Remote management only applies to MPP calls.
- 8 Close the System profile.

Configuring the Pipeline to interact with syslog

To maintain a permanent log of Pipeline system events and send Call Detail Reporting (CDR) reports to a host that can record and process them, configure the Pipeline to report events to a syslog host on the local IP network. Note that syslog reports are only sent out through the Ethernet interface.

To configure the Pipeline to send messages to a Syslog daemon:

- 1 Open the Ethernet > Mod Config menu.

```
90-C00 Mod Config
Log...
Syslog=Yes
Log Host=206.65.212.205
Log Port=514
Log Facility=Local0
```

- 2 Turn on Syslog.

- 3 Specify the IP address of the host running the Syslog daemon.

The host running a Syslog daemon is typically a UNIX host, but it may also be a Windows system. If the log host is not on the same subnet as the Pipeline, the Pipeline must have a route to that host, either via RIP or a static route.

Note: Do not configure the Pipeline to send reports to a syslog host that can only be reached by a dial-up connection. That would cause the Pipeline to dial the log host for every logged action, including hang ups.

- 4 Select Log Port and type the port number at which you want the Syslog host to listen for messages from this Pipeline.

The default port is port 514.

- 5 Set the log facility level.

This parameter is used to flag messages from the Pipeline. After you set a log facility number, you need to configure the Syslog daemon to write all messages containing that facility number to a particular log file. (That will be the Pipeline log file.)

- 6 Close the Ethernet profile.

To configure the Syslog daemon, you need to modify `/etc/syslog.conf` on the log host. This file specifies which action the daemon will perform when it receives messages from a particular log facility number (which represents the Pipeline). For example, if you set Log Facility to Local5 in the Pipeline, and you want to log its messages in `/var/log/Pipeline`, add this line to `/etc/syslog.conf`:

```
local5.info<tab>/var/log/Pipeline
```

Note: The Syslog daemon must reread `/etc/syslog.conf` after it has been changed.

Syslog messages

Syslog messages have a standard format that is described below. In addition to the normal traffic logged by Syslog, information may be generated for packets seen by the Secure Access firewall, if specified by SAM. By default, SAM will cause a syslog message to be generated for all packets blocked by a firewall.

Syslog messages use the format:

`<date> <time> <router name> ASCEND: <interface> <message>`

- `<date>` indicates the date the message was logged by syslog.
- `<time>` indicates the time the message was logged by syslog.
- `<router name>` indicates the router this message was sent from.
- `<interface>` is the name of the interface (ie0, wan0, and so on) or 'call' if the packet is logged by the call filter as it brings up the link.
- The `<message>` format has a number of fields, one or more of which may be present:

protocol The 4 hexadecimal digit Ether Type, or the network protocol name—"arp," "rarp," "ipx," "appletalk."
The protocol for IP protocols, is either the IP protocol number (up to 3 decimal digits) or one of the following names:

- ip-in-ip
- tcp
- icmp—In the special case of icmp, it will also include the ICMP Code and Type ([Code]/[Type]/icmp).
- udp
- esp
- ah

local	<p>For non-IP packets, is the source Ethernet MAC address of transmitted packets and the destination Ethernet MAC address of received packets. On a non-bridged WAN connection, the two MAC addresses will be all zeros.</p> <p>Local for IP protocols, is the IP source address of transmitted packets and the IP destination address of received packets. In the case of TCP or UDP, it will also include the TCP or UDP port number ([IP-address];[port]).</p>
direction	<p>An arrow “<-”, “->” showing the direction (receive and send respectively) in which the packet was traveling.</p>
remote	<p>For non-IP protocols, has the same format as “local” non-IP packets but shows the destination Ethernet MAC address of transmitted packets and the source Ethernet MAC address of received packets.</p> <p>For IP protocols, has the same format as <local> but shows the IP destination address of transmitted packets and the IP source address of received packets.</p>
length	<p>The length of the packet in octets (8-bit bytes).</p>
frag	<p>Used if the packet has a non-zero IP offset or the IP More-Fragments bit is set in the IP header.</p>
log	<p>Used to report one or more messages based upon the packet status or packet header flags. The packet status messages include:</p> <ul style="list-style-type: none">• corrupt, where the packet is internally inconsistent• unreachable, where the packet was generated by an “unreach=” rule in the firewall• !pass, where the packet was blocked by the data firewall• bringup, where the packet matches the call firewall• !bringup, where the packet did not match the call firewall• TCP flag bits that will be displayed include syn, fin, rst.• syn is only displayed for the initial packet which has the SYN flag and not the ACK flag set.
tag	<p>Contains any user defined tags specified in the filter template used by SAM.</p>

Using the Pipeline status windows

Eight status windows are displayed on the right side of the screen in the Pipeline configuration interface (Figure 8-1). These status windows provide a great deal of read-only information about what is currently happening in the Pipeline.

This section gives an overview of the information contained in the eight windows. Refer to the chapter entitled, “Status Windows Reference,” in the *Reference Guide* for a complete description of each line item in each status window.

10-100 1234567890 Link A NT1/CSU B1 * CARRIER B2	00-200 07:49:19 >M31 Line Ch LAN Session Up
20-100 Sessions > 1 Active ^ 0 corporate-gw 0	20-500 DYN Stat Qual N/A 00:00:00 OK 2 channels CLU 0% ALU 0%
20-300 WAN Stat >Rx Pkt: 72939069^ Tx Pkt: 64595101 CRC: 1350v	20-400 Ether Stat >RxPkt: 762800869^ Tx Pkt: 4595641 Col: 444314
00-100 Sys Option >Security Prof: 1 ^ Software +6.0b0+ S/N:5180736 v	00-400 HW Config >BRI Interface ^ Adrs: 00c08b43670 Enet I/F AUI v

Figure 8-1. Status windows

To scroll through the information in a status window, you must make the window active by hitting the TAB key until the window is highlighted by a thick border. If a lowercase v appears in the lower-right corner of a window, it means there is more information. You can see the additional lines by pressing the down-arrow key. For example, notice the lowercase v in the Sys Option window in Figure 8-1.

Performing system administration operations

This section describes the following system administration operations:

- Using DO commands to manually place and clear calls
- Restoring and saving a configuration
- Resetting the Pipeline
- Invoking the terminal server interface

Using DO commands

The DO menu is a context-sensitive list of commands that appears when you press Ctrl-D from the VT100 interface. The commands in the DO menu vary depending on the context in which you invoke it.

To initiate a DO command from the DO menu, press the number of the command. For example, press 1 to invoke the Dial command.

This is a complete list of DO commands:

- 0=ESC — Abort and exit the DO menu.
- 1=Dial — Dial the selected or current profile.
- 2=Hang Up— Hang up from a call in progress.
- 3=Answer — Answer an incoming call.
- C=Close Telnet — Close the current Telnet session.
- D=Diagnostics window
- E=Terminal Server
- P=Password — Log into or out of a Pipeline Security profile.
- S=Save — Save parameter values into the specified profile.

For details on each of these commands, see the *Reference Guide*.

To manually place a call, the Connection profile for that call must be open or selected in the list of Profiles. To clear a call, you can either open the Connection profile for the active connection, or tab over to the status window in which that connection is listed (see “Using the Pipeline status windows” on page 8-9).

To manually place a call:

- 1 Select or open the Connection profile for the destination you want to call.
- 2 Press Ctrl-D to invoke the DO menu.
- 3 Press 1 to invoke the Dial command.
- 4 Watch the information in Sessions status window. You should see the number being called followed by a message that the network session is up.

If you do not see the 1=Dial option, it may be because of these reasons:

- You are not in the correct profile.
- You do not have the appropriate security level enabled.
- You have not entered a dial number in the profile.
- You have not entered an IP address in the profile (if IP routing is enabled).

To manually clear a call:

- 1 Open the Connection profile or tab over to the status window that displays information about the active session you want to clear.
- 2 Press Ctrl-D to open the DO menu.

When you open the DO menu for an active session, it looks similar to this:

```
DO...
>0=ESC
  2=Hang Up
  P=Password
```

- 3 Press 2 to invoke the Hang Up command.

The status window will indicate when the call has been terminated.

Saving the Pipeline configuration

To save the Pipeline configuration using this method, you must have administrative privileges that include Field Service (such as the Full Access profile). And you must have a serial connection to the Pipeline.

Note: When you save the Pipeline configuration, the configuration data is written to a text file on the disk of the accessing host. *Passwords are not saved.*

Send and Recv passwords, Security profile passwords, and passwords specified in the Ethernet profile (Mod Config menu), are all set to the null password when you restore a configuration from a saved file. Be sure to record your passwords off-line if you need to restore them.

Before you start, verify that your terminal emulation program has a disk capture feature. Disk capture allows your emulator to capture to disk the ASCII characters it receives at its serial port. You should also verify that the data rate of your terminal emulation program is set to 9600 baud or lower and that the Term Rate parameter in the System profile (Sys Config menu) is also set to 9600. Higher speeds might cause capture errors.

You can cancel the backup process at any time by typing Ctrl-C.

To save the Pipeline configuration (except passwords) to disk:

- 1 Open the System > Sys Diag menu.

```
00-201 Restore Cfg
>00-202 Save Cfg
00-203 Sys Reset
00-204 Term Serv
```

- 2 Select Save Cfg and press Enter.

The following message appears:

```
Ready to download - type any key to start...
```

- 3 Turn on the Capture feature of your communications program and supply a filename for the saved profiles.

Consult the documentation for your communications program if you have any questions about how to turn on the Capture feature. Note that the HyperTerm and Terminal programs that ship with Microsoft Windows do not reliably save downloaded information.

- 4 Press any key to start saving your configured profiles.

Rows of configuration information are displayed on the screen as the file is downloaded to your hard disk. When the file has been downloaded to your hard disk, your communications program displays a message indicating the download is complete.

- 5 Turn off the Capture feature of your communications program.
- 6 Print a copy of your configured profiles for later reference.

If you examine the saved Pipeline data file, notice that some of the lines begin with START= and other lines begin with END=. These START/STOP lines and the block of data contained between them constitute a profile. If a parameter in a profile is set to its default value, it does not appear. In fact, you can have profiles with all parameters at their defaults and the corresponding START/STOP blocks would be empty. Make sure that there are no extra lines of text or characters either before START= or after END=. If there are, delete them; they could cause problems when you try to upload the file to the Pipeline.

The `tsave -a` command option supplies a listing of all parameter settings. To use `tsave -a`, you need access to a host with a TFTP server. To produce the listing, use Telnet to access the Pipeline unit. From the DO Command menu, select Diagnostics mode, and enter the command using the syntax shown below:

```
tsave -a nnn.nnn.nnn.nnn file.name
```

Where:

-a	Lists all the menu items in the software for the unit.
nnn.nnn.nnn.nnn	Is the local IP address of a host with a TFTP server.
file.name	Is the name of an empty file you create first in the TFTP boot directory of the host. Be sure you have read/write access to the file. (If you run into problems, the reason usually has to do with lack of read/write access.) The output file is written to the TFTP boot directory of the host.

Note: You can restore a configuration saved with `tsave -a` with the Diagnostics `trestore` command.

By default, the text configuration file you can create using the `tsave` command contains the VT100 interface parameter names. The `-m` option allows you to save the configuration file with the MIB field numbers instead.

To use the `tsave` command, you must use the diagnostic mode. From the DO menu, select D-Diagnostics. Then, to save the configuration of the Pipeline with the MIB field numbers instead of parameter names, enter this command line:

```
tsave -m <ipaddr> <filename>
```

For example:

```
tsave -m 200.253.164.100 all
```

This saves the entire configuration of the Pipeline with an IP address of 200.253.164.100 to a file called “all”.

Values are saved in the format:

OOOO:MMMM.FFFF

where

- OOOO represents the Occurrence number (if > 0),
- MMMM represents MIB Type (if > 0),
- FFFF represents the MIB field number (if MMMM > 0).

Note: You can restore a configuration saved with `tsave -m` with the Diagnostics `trestore` command.

Restoring the Pipeline configuration

To restore the Pipeline configuration, you must have administrative privileges that include Field Service (such as the Full Access profile, for example).

Before you start the restore procedure, verify that your terminal emulation program has an autotype (or ASCII file upload) feature. Autotype allows your emulator to transmit a text file over its serial port. You should also verify that the data rate of your terminal emulation program is set to 9600 baud or lower and that the Term Rate parameter in the System profile (Sys Config menu) is also set to 9600. Higher speeds might cause transmission errors.

You can use the `Restore Cfg` command to restore a full configuration that you saved by using the `Save Cfg` command, or to upload more specific configuration information obtained from Ascend, for example, a single filter stored in a special configuration file.

To load configuration information from disk:

- 1 Connect the backup device to the Pipeline Terminal port.
The backup device is typically the PC through which you access the VT100 interface.
- 2 Open the Sys Diag menu.
- 3 Select Restore Cfg and press Enter.
The following message appears:
`Waiting for upload data...`
- 4 Use the Send ASCII File feature of the communications software to send the Pipeline the configuration file.
If you have any questions about how to send an ASCII file, consult the documentation for your communications program. When the restore has been completed, the following message appears:
`Restore complete - type any key to return to menu`
- 5 Press any key to return to the configuration menus.
If you restored a complete configuration, the passwords used in your Security profiles have been wiped out. To reset the passwords:
- 6 Press Ctrl-D to invoke the DO menu, select Password, and choose the Full Access profile.
- 7 When you are prompted to enter the password, press Enter (the null password).
After you have restored your privileges by entering the null password, we recommend that you immediately open the Connection Profiles, Security Profiles, and Ethernet profile (Mod Config menu) and reset the passwords to their previous values.

See Appendix C, “Upgrading system software,” for related information.

Resetting the Pipeline

When you reset the Pipeline, the unit restarts and all active connections are terminated. All users are logged out and the default security level is reactivated. In addition, a system reset can cause a WAN line to temporarily be shut down due to momentary loss of signaling or framing information.

To reset the unit:

- 1** Open the Sys Diag menu.
- 2** Select Sys Reset and press Enter.
The Pipeline asks you to verify that you want to reset.

0=ESC
1=Reset
- 3** To confirm, type 1.

During a reset, the Pipeline clears active connections and runs its Power-On Self Test (POST), just as it would if the unit were power-cycled. If you do not see the POST display, press Ctrl-L.

While the yellow FAULT LED on the front panel is ON, the Pipeline checks its memory, configuration, installed modules, and lines. If any of the tests fail, the FAULT LED remains on or blinking.

The alarm relay remains closed while the POST is running and opens when the POST completes successfully. When you see this message:

```
Power-On Self Test PASSED  
Press any key...
```

Press any key to display the Main Edit Menu.

Using the terminal server interface

This section describes how to use the administrative commands that are available in the terminal server command-line interface.

Invoking and quitting the terminal server interface

To invoke the terminal server command-line interface, you must have administrative privileges. See “Activating administrative privileges” on page 8-3.

To open the command-line:

- 1 Open the Sys Diag > Term Serv menu and press Enter or, from the DO Command menu, select E=Termsrv.

The command-line prompt will be displayed at the bottom of the VT100 window:

```
ascend%
```

- 2 To close the command-line, use the Quit command at the prompt.

For example:

```
ascend% quit
```

The command-line interface closes and the cursor is returned to the VT100 menus.

Note: You could also use the Hangup or Local command to end the session. When a dial-in user enters the Local command, it begins a Telnet session to the Pipeline.

Terminal server commands

To display the list of terminal server commands, type:

```
ascend% ?
```

For help on a particular command, type that command followed by a question mark. For example:

```
show ?
```

Pipeline System Administration

Using the terminal server interface

The following table lists the terminal server commands, which are documented in detail in the “Terminal Server Commands” chapter of the *Reference Guide*.

Table 8-1. Terminal server commands

Command	Description
?	Displays help information.
dnstab edit	Starts editor for local DNS table.
dnstab entry	Displays local DNS table entry.
dnstab show	Displays local DNS table.
hangup	Closes the connection.
help	Help on any named command.
iproute	Displays information about IP routes in the unit’s routing table.
iproute add	Adds an IP route.
iproute delete	Deletes an IP route.
iproute show	Displays IP routes (same as show ip routes).
ipxping	Pings an IPX host.
local	Goes to local mode.
ping	Pings a remote host.
quit	Closes a terminal server session.
remote	Starts a remote management session.
set all	Displays current settings.
set arp clear	Clears ARP cache.
set fr	Frame Relay datalink control.

Table 8-1. Terminal server commands (continued)

Command	Description
set password	Enables dynamic password settings.
set sessid [val]	Sets and stores [val] or current ID.
set term	Sets the telnet/rlogin terminal type.
show arp	Displays the ARP cache.
show dhcp	Displays DHCP configuration parameters.
show dhcp address	Displays DHCP Address Assignment Information.
show dhcp lease	Displays DHCP lease Information.
show dnstab	Displays local DNS table.
show dnstab entry	Displays local DNS table entry.
show fr dlci [name]	Displays all DLCI information, or for [name].
show fr lmi	Displays Frame relay LMI information.
show fr stats	Displays Frame relay statistics information.
show icmp	Displays ICMP information.
show if stats	Displays interface statistics.
show if totals	Displays interface total counts.
show igmp clients	Displays IGMP clients.
show igmp groups	Displays IGMP groups table.
show igmp stats	Displays IGMP statistics.
show ip address	Displays IP address assignments.

Table 8-1. Terminal server commands (continued)

Command	Description
show ip routes	Displays IP routes.
show ip stats	Displays IP statistics.
show isdn	Displays ISDN events.
show netw networks	Displays NetWare IPX Networks.
show netw pings	Displays NetWare IPX Ping Stats.
show netw servers	Displays NetWare IPX Servers.
show netw stats	Displays NetWare IPX Statistics.
show revision	Displays system revision.
show sessid	Displays current and base session ID.
show tcp connection	Displays TCP connection table.
show tcp stats	Displays TCP statistics.
show udp listen	Displays UDP listen table.
show udp stats	Displays UDP statistics
show uptime	Displays system uptime.
tcp	Opens a raw TCP/IP session to an IP host.
telnet	Establishes a telnet session with another host.
test	Tests your ISDN line by calling itself.
traceroute	Lets you trace a route to a host.

Accessing a local Pipeline via Telnet

If a remote user Telnets to the Pipeline and the Ethernet > Mod Config > Telnet PW has been set, the user is prompted for the Telnet password. Local users Telnetting to the Pipeline over the Ethernet must also supply this password.

The Telnet password verification trap reports the IP address of the Telnet client whose login attempts failed. The address is included in the security violation message issued whenever the maximum number of Telnet login attempts to a Pipeline has been exceeded.

To Telnet into a Pipeline, a user must supply the appropriate password, which is then verified. If the user cannot supply the correct password, an SNMP trap message is sent to all SNMP clients enabled for SNMP security messages.

The message includes the following information:

- The session number for the attempted Telnet session.
- The IP address of the host (the Pipeline).
- The associated IP address of the Telnet client that attempted the connection.

The format of the message is as follows:

```
mm.mmm.mmm.mmm  Enterprise Specific Trap (15) Uptime: xx:xx:xx
Name.iso.org.dod.internet.private.enterprises.ascend.sessionStatus Group.
IpAddress: ttt.ttt.ttt.ttt
sessionStatusTable.sessionStatusEntry.ssnStatusUserIPAddress%d
```

Where:

mm.mmm.mmm.mmm	Host's IP address
ttt.ttt.ttt.ttt	Telnet client's IP address
%d	attempted Telnet session number

APP Server utility

This appendix includes these topics:

About the APP Server utility	A-1
APP Server installation and setup	A-2

About the APP Server utility

The Ascend Password Protocol (APP) Server utility lets you respond to token password challenges received from an external network authentication server (NAS). These external authentication servers typically change passwords many times a day, and sync up with hand-held personal security cards to provide users with the current password in real-time. The LCD on the users' security cards displays the current password required to gain access at that moment to the secure network.

Whenever you require a connection to a secure network, the Pipeline initiates the call and negotiates an initial session. The NAS returns a password challenge, which the Pipeline passes to the APP Server. Once you answer the challenge correctly, you are connected to the secure server.

To make this happen, obtain a copy of the APP Server utility from the Ascend FTP site and install it on your computer. The steps to do this are listed in "APP Server installation and setup" on page A-2. After the installation, each time your computer boots, the APP Server starts and runs in the background.

Then configure your Pipeline to communicate with the APP Server utility. The steps to do this are listed in “Configuring the Pipeline to use the APP server” on page A-2.

APP Server installation and setup

The APP Server utility is provided for Macintosh, DOS, Windows 3.1, Windows 95, Windows NT, and UNIX. The utility is available from the Ascend FTP server. The files can be found at ftp.ascend.com/pub/Software-Releases/AppServer. From this location, select the folder for your operating platform and download the self-extracting archive.

Configuring the Pipeline to use the APP server

APP is a UDP protocol whose default port is 7001. The communication between the Pipeline and the host running the APP Server may be unicast (when both the Pipeline and the host have an IP address) or broadcast (when the host may not have an IP address).

To set up the Pipeline to communicate with the APP Server utility, do the following:

- 1** Open the Ethernet > Auth profile.
- 2** Set the APP Server parameter to Yes.

`APP Server=Yes`

This enables the Pipeline to communicate password challenges to the host running the APP Server utility.

- 3** Specify the IP address of the host (that is, the computer) running the APP Server utility.

For example:

`APP Host=10.65.212.1`

If the host obtains its address at boot time from a BOOTP or DHCP server, or if it has no IP address, you can specify the IP broadcast address in this parameter (255.255.255.255).

- 4** Specify the UDP port to use for communicating with the host running the APP Server.

For example:

APP Host=7001

7001 is the default UDP port for the APP Server.

Note: If you change this number, you must specify the new UDP port number in the Password AppServer Control Panel (Macintosh), APP Server utility (DOS), the WIN.INI file (Windows), or /etc/services (UNIX). The Pipeline and the host running the APP Server utility must agree on the UDP port number.

- 5 Close the Ethernet profile.

Using App Server with Axent SecureNet

When using SecureNet, you must install a Softkey on your computer's hard drive, or supply a diskette-based Softkey that needs to be inserted in your computer's floppy drive when logging onto a SecureNet system. If the Softkey is present when App Server is installed, the App Server INI file (or Password AppServer Control Panel file on a Mac) is automatically modified to work with the Axent SecureNet Softkey. (If the Softkey is installed after App Server, you can manually modify the Path key in the WinSNK section of the INI file, as shown below.)

The App Server functions as usual with Softkey, except that whenever App Server is started it attempts to find the Softkey, and if found, the Axent SecureNet software prompts for a PIN. Once entered, all subsequent transactions between the authentication server and the App Server are transparent, unless an error occurs, or if the Softkey has expired.

Creating banner text for the password prompt

You can create a banner that greets users when a challenge message is received. The APPSRVR.INI file, in the directory in which the APP Server utility is installed, should contain banner text to be displayed along with the password prompt when a challenge message is received. The banner can be up to 200 characters and up to five lines of text. To set up the BANNER on a Macintosh, use the information below and enter it in the Password AppServer Control Panel. Also see "Installing APP Server on a Macintosh" on page A-13.

In the APPSRVR.INI file, the first line of the file must contain the text “[BANNER]”.

For example:

```
[ BANNER ]  
line1=The security password has changed. Please consult your  
line2=security card and enter the current password now.  
line3=You have 60 seconds to enter the new password.
```

The banner is followed by the challenge prompt in the APP Server screen. A user has 60 seconds to obtain the current password from the security card and enter it correctly.

There are three sections in AppSrvr.ini. The sections are described in the following table:

Table A -1. APP Server INI file contents

INI section	Description
[BANNER]	Up to 5 line of text, each one must begin with the syntax “line x=”, where x is a number from 1 to 5. For example [BANNER] line 1=“First line of text” line 2=“Second line of text” ...
[PROFILE]	Allows for the following two key names: Name = User = Name is the name of the remote Ascend unit. Note: This field is ignored when using Axent SecureNet since this information is contained in the Softkey authentication routine.) User is the profile name to use when connection.

Table A-1. APP Server INI file contents (continued)

INI section	Description
[WinSNK]	<p>Consists of 33 lines with the first using the key name, Path, and all remaining lines using a number from 0 to 31.</p> <p>Path is the fully qualified path to the location of the installed Axent SecureNet Softkey. The purpose of this section is to maintain a list of text messages received from the authentication server, which allows you to keep App Server synchronized with any change made by the SecureNet administrator.</p> <p>0-31 contain the text as entered on the authentication server.</p>

Additionally, a section entitled [App Server] is added to WIN.INI when App Server is installed, containing the default socket data (automatically entered by the App Server utility). Even though the data is listed in WIN.INI, the values are actually stored in the Windows Registry.

Two keys are included in the [App Server] section of WIN.INI:

- udp_port
- bcast_udp_port

The following is a sample AppSrvr.ini file that illustrates the overall format.

```
[BANNER]
line1="This is a sample."
[PROFILE]
Name=hummer
User=administrator
[WinSNK]
Path=F:\WinSNK
0=Call intercepted by Defender Security Server
1=Unauthorized use of this system is prohibited
3=Enter ID:
4=SNK Challenge: %s ^M^JEnter Response:
5=Invalid Identification.
6=Invalid SNK Response^M^JSNK Challenge: %s ^M^JEnter
Response:
7=Access Approved. You are now connected to service.^M^J
```

```
8=Access Denied.^M^J
9=All Channels of Security Server are busy. Try again later
^M^J
10=Unexpected packet from Agent^M^J
11=Cannot start new call on active channel^M^J
12=Cannot start new call on active channel^M^J
13=Unexpected input from user.^M^J
14=Enter Password:
15=Invalid Identification.^M^JEnter ID:
16=Your password has expired.^M^JEnter New Password:
17=Enter New Password:
18=Enter New Password again:
19=Passwords didn't match.^M^JEnter New Password:
20=Outside your time class.^M^J
21=Outside your date class.^M^J
22=New password must differ from old.^M^JEnter New Password:
23=New password is too short.^M^JEnter New Password:
24=New password must include numeric digit.^M^JEnter New
Password:
25=Request noted.^M^JEnter old password
26=Your account is locked due to excess violations.^M^J
27=Your ID is already active on another channel.^M^J
28=Your password has been changed.^M^J
29=Your account is locked due to non-usage.^M^J
30=You are not authorized for that host.^M^J
31=Inactivity Timeout.^M^J
```

Installing and using the UNIX APP Server

When a user starts an application that requires a connection to a host on a secure network, the Pipeline initiates the call as usual. After the initial session negotiation, the remote ACE or Safeword server returns a password challenge that looks similar to this:

```
From: hostname
0-Challenge: challenge (or null challenge, depend-
ing on your setup)
Enter next password:
```

This prompt is displayed in the APP Server screen on the UNIX host. A user has 60 seconds to obtain the current dynamic password from the security card and

enter it correctly. If multiple users need to use the APP Server, the user can include a name in this format:

`password.username`

(A password followed by a period, followed by the user name.)

To install the APP Server utility on a UNIX host:

- 1** Edit the Makefile appropriately for your operating system and compiler.
- 2** Compile the `appsvr` source file (`make`).
- 3** Add a line to `/etc/services` assigning UDP port 7001 to the APP Server utility.

If you can use the default UDP port 7001 (if it is not already assigned), add this line to the `/etc/services` file to document that the port is now in use:

`appServer<tab>7001/udp`

If port 7001 is already assigned to a different application, you can use a different port for the APP Server utility by adding a line such as this to the services file:

`appServer<tab>nnn/udp`

where *nnn* is the port number to be used. Make sure that the Pipeline configuration agrees with this number.

- 4** If the UNIX host has an IP address, you can run the utility in unicast mode by typing this command at the UNIX prompt:

`./appsvr`

When you run the utility in unicast mode, it transmits packets on the specified UDP port with the source address set to its own IP address. When the Pipeline receives those packets on the specified UDP port, it returns packets to that IP address.

- 5** If the UNIX host does *not* have an IP address (for example, if it obtains its address from a BOOTP or DHCP server), run the utility in broadcast mode instead by typing this command:

`./appsvr -b`

The `-b` option sets a socket option to allow broadcast transmissions and inhibits the utility's complaints about receiving invalid APP frame types when it receives its own transmissions.

Note: On some UNIX systems, you need root privileges to run the APP Server utility in broadcast mode. (Some hosts disallow broadcast transmissions without root privileges.) If you are running the utility in broadcast mode, make sure that the Pipeline is configured with the broadcast address in the APP Host parameter (APP Host=255.255.255.255).

Installing and using the APP Server utility for DOS

To initiate a connection to a remote secure network, the DOS user reboots the PC. After the initial session negotiation, the remote ACE or SAFEWORD server returns a password challenge that looks similar to this:

```
From: hostname
0-Challenge: challenge (or null challenge, depend-
ing on your setup)
Enter next password:
```

If more than one user uses the APP Server to log into a remote secure network through the Pipeline, each user must include a user name in this format:

```
password.username
```

The syntax is a password followed by a period, followed by the user name.

The DOS version of the APP Server utility requires an ODI driver for its networking needs. It must be installed in AUTOEXEC.BAT immediately after loading the ODI driver. (You may need to edit STARTNET.BAT to accomplish this; however, this version no longer requires any changes to NET.CFG.)

To install the APP Server utility for DOS:

- 1 Create an \ASCEND directory below the root directory.
- 2 Copy APPSRVDS.EXE into that directory.
- 3 If the APPSRVR.INI exists, copy that into the directory as well.
See "Creating banner text for the password prompt" on page A-3.
- 4 Open AUTOEXEC.BAT and add a command line invoking APPSRVDS.EXE.

The APPSRVDS.EXE DOS utility does not require an IP stack or IP address, but it does require an ODI driver.

The command line for APPSRVDS.EXE must be positioned after the line invoking the network ODI driver and *before* the network protocol stack (TCP/IP or IPX or other supported protocol). For example:

```
C:\NOVELL\LSL.COM
C:\NOVELL\XXXODI.COM
C:\ASCEND\APPSRVDS.EXE

REM Protocol Stack is loaded next
```

5 Close AUTOEXEC.BAT.

6 Reboot.

There are several options you can use in the AUTOEXEC.BAT command-line:

- /t — specifies a time delay between connection attempts (sec)
- /y — specifies the number of cycle counts (attempts to connect) before timeout
- /m — specifies the MAC address (in decimal) of the PC running the utility
- /p — specifies a UDP port number for communicating with the Pipeline
- /b — specifies a UDP port for broadcast message
- /f — suppresses the call at startup
- /d — disconnects the call
- /c — specifies the name of the Connection profile to use to connect to the remote secure network
- /? — displays a help screen

Note: The PC sends a broadcast UDP packet that has the destination and the source port 7001 unless you specify otherwise with the /p or /b options. If you specify a number other than 7001 in the APP Port parameter, you must use one of these options to specify the same port.

If no command-line variables are specified, the APP Server utility uses the following default values:

- Time delay between connection attempts = 20 seconds
- Number of cycles is set to 3 (3 times 20 seconds)
- APP Server PC MAC address = none (zeros)

APP Server utility

APP Server installation and setup

- UDP port to use = 7001
- Broadcast UDP port is the same as communication UDP port
- APP Server will force a connection upon execution

Note: A Connection profile defined in the Pipeline is required to log into the remote secure network, so if the APP Server line in AUTOEXEC.BAT does not specify Connection profile name, the user will be prompted for one as the system boots.

For example, this command:

```
C:\ASCEND\APPSRVDS.EXE /Chicago /t20 /p7005
```

specifies a Connection profile named “Chicago,” assigns a 20-second time delay between connection attempts, and designates UDP port 7005 for communicating with the Pipeline.

```
C:\ASCEND\appsrvds.exe /Chicago /m00805110C7A44 /  
p7523 /t65 /b7112
```

specifies a Connection profile named “Chicago,” specifies 00805110C7A44 as the MAC address of the PC running the utility, designates UDP port 7523 for communicating with the Pipeline, assigns a 65-second time delay between connection attempts, and designates port 7112 for sending broadcast messages to initiate a call.

Installing and using the APP Server utility for Windows

The user interface is the same for all Windows versions of the APP Server utility, although the utility itself and the way in which it is installed differs.

To use the Windows utility:

- 1 If the utility is not already running, start it by using the Services applet on the Control Panel.

- 2 Click Connect.
A Settings dialog box opens (shown below).
- 3 Enter the name of the Connection profile used to log into the remote secure network.
- 4 Enter your user name.
The name you enter must be no longer than 32 characters and cannot contain spaces. Once entered, it is saved to disk and appears as the default the next time you log on.
- 5 Click OK.
After the initial session negotiation, the remote ACE or SAFEWORDD server returns a password challenge, which is displayed in its own dialog box. A user has 60 seconds to obtain the current dynamic password from the security card and enter it correctly.
- 6 Type the current password and click OK.
- 7 To log out of the remote network, click Disconnect.
- 8 Type the name of the Connection profile that defines your connection to the remote network, and then click OK. Once entered, it is saved to disk and appears as the default the next time you log on.

Installing the APP Server utility for Windows 3.1

To install the APP Server on a Windows 3.1 system:

- 1 Create an \Ascend directory below the root directory.
- 2 Copy APPSRV31.EXE into that directory.
- 3 If the APPSRVR.INI exists, copy that into the directory as well.
See “Creating banner text for the password prompt” on page A-3.
- 4 Copy CTL3D.DLL into the Windows System directory.

We recommend adding the APP Server utility to the startup group (provided that the network, including WINSOCK, is started as part of normal system startup.

To create an icon and add the APP Server to the startup group:

- 1 Create a new program group in your Program Manager.
Choose File > New > Program Group, and type **Ascend**.
- 2 Create an icon for APPSRV31.EXE in your Program Manager.

Choose File > New > Program Item.

- 3 To launch the APP Server utility when you start Windows, place the APPSRV31.EXE icon in your Startup group.
If you prefer not to add the APP Server utility to your Startup group, you can launch the utility manually by double-clicking its icon.
- 4 Reboot.

Installing the APP Server utility for Windows 95

To install the APP Server on a Windows 95 system:

- 1 Copy the file XAS-W95.EXE into a temporary directory.
XAS-W95.EXE is a self-extracting zip file.
- 2 Execute the file from the DOS shell.
It will expand to several files which include the Windows 95 Setup program.
- 3 From the START menu, run the Setup program in this directory.
- 4 Follow prompts and select the destination directory where the APP Server for Windows 95 should be installed.

The APP Server for Windows 95 will start automatically whenever the system reboots. You may close the APP Server in a session, but next time the system is rebooted, it will start again.

To permanently remove or disable the APP Server, you must edit the Windows 95 Registry to remove the key that references APPSRV95.EXE.

Installing the APP Server utility for Windows NT

To install the APP Server on a Windows NT system:

- 1 Copy the file XAS-NT.EXE into a temporary directory.
XAS-NT.EXE is a self-extracting zip file.
- 2 Execute the file from the DOS shell.
It will expand to several files which include the Windows NT Setup program.
- 3 Run the Setup program in this directory.

- 4 Follow prompts and select the destination directory where the APP Server for Windows NT should be installed.

The APP Server for Windows NT will start automatically whenever the system reboots. You may close the APP Server in a session, but next time the system is rebooted, it will start again.

There are three icons provided during installation which enable you to temporarily disable the APP Server, manually control when it runs, or remove it from the system.

- **Activate service icon**
Running the activate service icon will stop the service if it is running and then restart or activate it.
- **Remove service icon**
Running the remove service icon will stop the service if it is running and remove it from the service database; it will no longer be listed as a service by the Services applet on the Control Panel.
- **Uninstall service icon**
Running the uninstall service icon will cause the files, icons, program groups, and registry entries to be removed from the system.

Installing APP Server on a Macintosh

Execute the file Install Password AppServer to install the software. Easy install is selected by default; just click Install to complete the installation and start the Password AppServer. The Password AppServer automatically starts up each time the system is booted.

Open Transport is required for proper operation of the Password AppServer for Macintosh.

Configure your Pipeline as described in “Configuring the Pipeline to use the APP server” on page A-2.

To use BANNER, start the Control Panel named Password AppControl and enter the desired text for each line. Note that five lines or less may be entered. Each line may contain text or be blank. The text entered here will be displayed along with the password prompt.

Troubleshooting

B

This appendix includes the following topics:

Cabling problems: Rule these out first	B-1
Common problems and their solutions	B-2
Problems configuring the Pipeline	B-5
ISDN BRI interface problems	B-8
Problems accessing the remote network	B-13

Cabling problems: Rule these out first

If you're unable to establish a connection with a remote network, first check that the ISDN line is plugged into the Pipeline. Telephone companies report that this is the most common cause of initial failures.

Another common problem is incorrect Ethernet cabling. The cross-over cable provided in the Pipeline package can be used only in a direct connection between the Ethernet adapter (or external transceiver) in the computer and the Pipeline. If you are connecting the Pipeline to a 10BaseT hub, you must use a regular 10BaseT cable between the hub and the Pipeline, and between the hub and the computer.

For Macintosh computers, sometimes the port you used to plug the serial cable into the Macintosh doesn't work. You can use either the modem or printer port in the Macintosh. If one doesn't work, try the other one.

See “Check the installation” on page B-13 for related information.

Common problems and their solutions

This section lists problems you might encounter and describes ways to resolve them.

General problems

When the list of DO commands appears, most operations are not available

You might need to select a specific Connection profile in order to see certain DO commands. For example, to dial a Connection profile, you must move to the Connection profile in the Connections menu, and then type Ctrl-D 1.

Note that you cannot dial if Operations=No for the control port. If a call is already active, DO 2 (Hang Up) appears instead of DO 1 (Dial).

If you do not see the DO 1 (Dial) option, it may be because:

- You are not in the correct profile.
- You do not have the appropriate security level enabled.
- You have not entered a dial number in the profile.
- You have not entered an IP address in the profile (if IP routing is enabled).

Profile configuration problems

The most common problems result from improperly configured profiles.

The T1 line is in service, but the Pipeline cannot dial a call

Verify that the Connection profile is correctly configured by following these steps:

- 1** Make certain that you have entered the correct phone number to dial.
- 2** Check that the Data Svc parameter specifies a WAN service available on your line.

If you request a WAN service that is not available on your line, the WAN rejects your request to place a call.

- 3** Determine whether you have correctly set the parameters controlling Dynamic Bandwidth Allocation.

For detailed information, see Chapter 1, “Configuring WAN Connections.”

The data appears to be corrupted on 1 Chnl or 2 Chnl calls dialed in the U.S. to another country

On some international calls, the data service per channel is not conveyed by the WAN to the Pipeline answering the call. You must therefore set Force 56=Yes in the Connection profile. If you do not, the Pipeline incorrectly thinks that the call uses 64-kbps channels.

The first channel of an MP+ call connects, but then the call clears or does not connect on the remaining channels

The most common error in defining Connection profiles is specifying incorrect phone numbers. The Pipeline cannot successfully build inverse multiplexing or MP+ calls if the phone numbers in the Connection profile of the called unit are incorrect. The phone numbers that you specify in the Connection profile are the numbers local to your unit. Do not enter the phone numbers of the Pipeline you are calling in the Connection profile.

When the Pipeline tries to place a call, the error message No Channel Avail appears in the Message Log display

Check the configuration of your line in the Configure profile.

This message can also indicate that the T1 cables have been disconnected or were installed incorrectly.

Hardware configuration problems

If you cannot communicate with the Pipeline through the VT-100 control terminal, you might have a terminal configuration, control port cable, or Pipeline hardware problem.

No data is displayed on the VT-100

If the Pipeline is in this state, verify that the unit completes all of the power-on self tests successfully by following these steps:

- 1 Verify that the Pipeline and your terminal are set at the same speed.
- 2 Locate the LED labeled CON.
- 3 Switch on the Pipeline.

The CON LED should remain off except during the power-on self tests. If you are using the Control Monitor, type Ctrl-L to refresh the screen.

If the CON LED remains on longer than a minute, there is a Pipeline hardware failure. A blinking CON LED also indicates a hardware failure.

Should these situations arise, contact Ascend Communications, Inc. Customer Support.

The CON LED is off, but no data is displayed on the Control Monitor's VT-100 terminal

If the unit passed its power-on self tests and you still cannot communicate with the Control Monitor, type Ctrl-L to refresh the screen. If you still do not see any data, check the cabling between the Pipeline and your terminal by following these steps:

- 1 Check the pin-out carefully on the 9-pin cable.
The control terminal plugs into the HHT-VT-100 cable or 9-pin connector labeled Terminal on the back of the Pipeline. If you are connecting to an IBM PC-like 9-pin serial connector, a straight-through cable is appropriate. Otherwise, you might need a 9-to-25 pin conversion cable.
- 2 Check the flow control settings on your VT-100 terminal.
If you are not communicating at all with the Pipeline, see whether you can establish communications after you have turned off all transmit and receive flow control at your terminal or terminal emulator.
- 3 Determine whether you need a null-modem cable converter.
In general, these are not required for communications to the Pipeline. However, so many different cable and terminal configurations are available that occasionally a null-modem cable converter might be required.

Random characters appear on the Control Monitor screen

If random or illegible characters appear on your display, there is probably a communications settings problem. You must make these settings:

- 9600 bits per second data rate
- 8 data bits
- 1 stop bit
- No flow control
- No parity

If you have changed the data rate through the Sys Config menu, make certain that your VT-100 terminal matches that rate.

Also, make sure the Term Rate setting matches the speed configured for your Com Port.

Use Ctrl-L to refresh the screen.

The start-up display indicates a power-on self test failure

If the start-up display indicates a failure in any of its tests, an internal hardware failure has occurred with the unit. In this case, contact Ascend Communications, Inc. Customer Support.

Problems configuring the Pipeline

There are two common problems associated with the Pipeline configuration procedure:

- The communications program does not display a profile when you press Ctrl-L.
- A profile appears when you press Ctrl-L, but it isn't the Configure profile shown in this manual.

If you see garbage characters on the screen, make sure that vt100 emulation is set to the right speed (9600 bps).

No profile appears in your communications program

If no profile appears when you press Ctrl-L in your communications program, one of these conditions could be causing the problem:

- Your Pipeline is not receiving power.
- Your Pipeline is not connected to the serial port of your computer.
- Your communications program is not configured correctly for your Pipeline, or it is not communicating on the right port.
- There is a hardware problem with the Pipeline.

To diagnose and solve the problem, follow these steps:

- 1 Check the pwr LED on the front panel of the Pipeline.

If the pwr LED is not on, the unit is not receiving power. It may not be connected to a power source. Continue to step 2.

If the light is on, continue to step 4.

- 2 Connect your Pipeline to a power source.

If your Pipeline is plugged into a power strip or surge protector, make sure the power strip or surge protector is plugged in and turned on.

Once you are sure the Pipeline is connected to a power source, if the pwr LED is on, continue to step 3.

If the pwr LED is still not on, contact the Ascend Technical Assistance Center at 1-800-ASCEND-4.

- 3 Check the con LED.

If the con LED goes off within thirty seconds after you connect the Pipeline to a power source, continue to step 4.

If the con LED is blinking or on more than thirty seconds after you have connected the Pipeline to a power source, contact the Ascend Technical Assistance Center at 1-800-ASCEND-4.

- 4 Press Ctrl-L to refresh the screen.

If no profile appears, continue to step 5.

If a profile appears, but it isn't the Configure profile, go to "A profile appears but it isn't the Configure profile" on page B-8."

- 5 Check to see if your Pipeline is connected to your computer's serial port.

If necessary, connect the Pipeline to your computer and continue to the next step.

If your Pipeline is connected to your computer, continue to step 6.

- 6** Press Ctrl-L to refresh the screen.

If no profile appears, continue to step 7.

If a profile appears, but it isn't the Configure profile, go to "A profile appears but it isn't the Configure profile" on page B-8."

- 7** Check to see if your communications program is configured for the Pipeline. Your communications program should be configured as follows:

- VT100
- 9600 bits per second
- 8 data bits
- No parity
- 1 stop bit
- No flow control
- Direct connect

If necessary, configure your communications program, then continue to the next step.

- 8** Press Ctrl-L to refresh the screen.

If no profile appears, contact your network administrator.

If a profile appears but it isn't the Configure profile, continue to the next section.

A profile appears but it isn't the Configure profile

If a profile appears, but it isn't the Configure profile, your Pipeline may already have been configured.

Solving this problem is easy: press Escape until you reach the Main Edit Menu, and then select Configure.

ISDN BRI interface problems

Provisioning or switch type problems

If voice calls are not being received correctly, it's possible that your ISDN line was provisioned incorrectly at the central office switch.

If you are unable to receive a voice call while a data call is in progress, it's possible that your line was configured with the Point-to-Point switch type. In cases where both B channels are in use for a multi-channel data call, the Point-to-Point switch is not able to pass on a voice call for the Pipeline to service.

If you suspect a provisioning or switch type problem, call the telephone company and work through the provisioning information described at the beginning of this guide.

SPID format problems

If the SPIDs entered in the Pipeline configuration are incorrect, the Pipeline will be unable to access the ISDN line.

The most common problem with SPIDs is that they were entered incorrectly, either by mistake or because the telephone company provided the wrong information. If wrong or incomplete information was provided about the SPID numbers assigned to your ISDN line, try adding 00 to the end of the SPID number. Or, if the suffix ends in a double digit, such as 01 or 02, try replacing those two digits with a single digit, such as 1 or 2. If neither of these suggestions works, call the telephone company and request that they verify the SPIDs you have.

Dialing and answering do not operate reliably

To resolve this problem, follow these steps:

1 Check your cabling.

The first and most critical aspect of ISDN BRI interfaces is the cable or cables connecting the Pipeline to the WAN line or WAN-terminating equipment. Typically, WAN interface cabling problems appear immediately after installation. If you are unsure about the cabling required for your application, contact Customer Support. See the Specifications appendix in the *Start Here* booklet. It describes the general ISDN BRI interface requirements and lists cabling pin-outs.

2 For a T1 line, if the cabling is not the problem, check that the value of the Buildout parameter in the Nailed T1 Group profile matches the actual distance in your configuration.

- Contact your carrier representative to determine which value to choose.

The status of an ISDN BRI line in the WAN Status windows is No Logical Link

In some countries outside the U.S., it is common for no logical link to exist before the Pipeline places a call.

In the U.S., when you first plug a line into the Pipeline or switch power on, the central office switch can take as long as 15 minutes to recognize that the line is now available. You might have to wait that long for the line state to change to Line Active (LA). The physical link can exist without a logical link up on the line.

If you wait longer than 15 minutes and the line is still not available, follow these steps:

1 Check whether all the ISDN telephone cables are wired straight through.

If you are running multipoint (passive bus) on your switch, all of the ISDN telephone cables must be wired straight through. If any of the cables are wired to cross over, you will not be able to place calls.

2 Check that 100% termination is provided on each ISDN line.

3 Check whether you have correctly specified the SPIDs (Service Profile Identifiers) in the Configure profile for each line.

If the SPIDs are not correctly specified, the line status might indicate No Logical Link. Check with your system manager or carrier representative to obtain the SPID or SPIDs for your line. You specify your SPIDs in the Configure profile.

T1 and ISDN BRI circuit-quality problems

The Line Status window indicates that the T1 line is in a Red Alarm state

If the line is in a Red Alarm state, the Pipeline cannot establish proper synchronization and frame alignment with the WAN. This behavior is normal for as long as 30 seconds when a T1 line is first plugged into the Pipeline.

If the Red Alarm condition persists for longer than 30 seconds, follow these steps:

- 1** Check the value of the Framing Mode parameter in the Nailed T1 Group profile.
Change the value to the other available option and check to see whether the Red Alarm condition goes away within 30 seconds.
- 2** If the Red Alarm state still persists, check the cabling.
You might have a crossover cable installed when a straight-through cable is required, or vice versa. If the Pipeline is connected through bantam plugs, reverse the transmit and receive plugs. Then, allow the Pipeline to attempt to establish synchronization for an additional 30 seconds.
- 3** You can eliminate the T1 cabling as a possible cause by replacing the T1 connection with a loopback plug.

A T1 line is in use and the Line Status window indicates an ALARM state

An ALARM state means that the physical configuration of the T1 line is correct, but that the D channel is not communicating with the WAN. To resolve this problem, follow these steps:

- 1** Verify with your carrier representative that the D channel is channel 24, or the number configured in the Configure profile.
- 2** If the channel number is correct, check the value of the Line Encoding parameter in the Nailed T1 Group profile.

When B8ZS encoding is in use, a non-inverted D channel is established. If AMI encoding is selected, an inverted D channel is established. Check the line translations provided by your carrier representative and set the line encoding to match the inversion requirements.

- 3 Check whether you have specified the proper Buildout value in the Nailed T1 Group profile.

- 4 Check whether the D channel is in service.

If no equipment has been plugged into the line for a short period of time, namely five to ten minutes, the D channel is taken out of service. You might need to contact your carrier to put the D channel back into service.

T1 access problems

The Pipeline never uses some T1 channels

To resolve this problem, follow these steps:

- 1 Check whether the T1 line has been recently connected to a device that does not support the full 24 channels.

If such has been the case, the switch might take the unused channels out of service. This situation can arise on either the local or the remote end.

- 2 Check whether the channels enabled in your Nailed T1 Group profile correspond to the channels enabled in the circuit.

If only some of the channels in the circuit are available for data calls, you must specifically enable those channels in your Nailed T1 Group profile.

- 3 If you place a call and some channels are always skipped, call your carrier representative.

An outgoing call that uses a T1 access line fails to connect to the remote end

If the Connection profile is configured correctly and you still cannot place an outgoing call, check the service state of the T1 access line.

Frequently, if a T1 access line has been unplugged for an extended duration, the switched services available on the line are taken out of service. Once you install the Pipeline, you might need to call your carrier representative to have the line

reactivated. If this is the case, leave the Pipeline on all the time, even when you are not using it; otherwise, you will have to call your carrier representative to reactivate the line each time the unit is switched off and on.

Bridge/router problems

The quality of the link is questionable

When running FTP (File Transfer Protocol), the data transfer rate appears in bytes per second. Multiply this rate times 8 to get the bits per second. For example, suppose that you are connected to Detroit on a 56-kbps B channel and that FTP indicates a 5.8 Kbps data rate; in this case, the link is running at $5.8 \times 8 = 46.8$ kbps, or approximately 83% efficiency. Many factors can affect efficiency, including the load on the FTP server, the round-trip delay, the overall traffic between endpoints, and the link quality.

You can check link quality in the WAN Stat status menu, or by running a ping between the same endpoints. Dropped packets hurt the link's efficiency, as does round-trip delay. Random round-trip delay indicates heavy traffic, a condition that also drops the efficiency of the link.

The Pipeline hangs up after answering an IP call

To resolve this problem, follow these steps:

- 1** If you are running PPP, check that you have entered the proper passwords.
- 2** Check that Auth is set to PAP or CHAP.
- 3** If you are routing IP over PPP, check that the calling device gives its IP address.

Some calling devices supply their names, but not their IP addresses.

However, you can derive an IP address if the calling device is listed in a local Connection profile. Try enabling PAP or CHAP for the Recv Auth parameter so that the Pipeline matches the caller's name to the Station parameter in a Connection profile and gets the corresponding LAN Adrs.

Problems accessing the remote network

If, when you press Ctrl-D in the Configure profile, the status window in the upper right corner displays a message other than LAN Session Up, you should first disconnect the Pipeline from the phone line connection, reconnect it, then try accessing the remote network again. If you still cannot access the remote network, one or a combination of the following may be a problem:

- Your Pipeline may not be installed correctly.
- Your Pipeline may not be configured correctly.
- Your phone line may not have been activated, or there may be a problem with the telephone network.

Check the installation

- 1** Make sure your Pipeline is connected to your phone line.
- 2** Check the WAN LED on the front panel of your Pipeline.
If the WAN LED is not blinking, continue to the next section, “Configuration problems” on page B-14.
If the WAN LED is blinking, one of the following may be the case:
 - Your Pipeline may not be connected to the phone line.
 - If you do not have an integrated NT1 interface, your Pipeline may not be connected to an NT1.
 - Your phone line may not be activated.
 - Your ISDN channel may be temporarily unavailable.
 - You may have entered an incorrect switch type. Check the setting in the Configure profile.
 - You may have entered the wrong SPID. Check the setting in the Configure profile, and confirm the values with your service provider.
- 3** Check to make sure you have connected your Pipeline to your ISDN line.
If necessary, connect your Pipeline to your ISDN line. If your Pipeline does not have an integrated NT-1 interface, make sure it is connected to an NT-1, and that the NT-1 is connected to the ISDN line as shown in your NT-1 manual.

Once you are connected, if the WAN LED is still blinking, continue to step 4.

- 4 Contact your ISDN service provider to see if your lines have been activated. If they have been activated, check to see if your service provider is experiencing problems with their telephone network.

If your lines are not activated, wait until they are, then try the call again.

If your service provider is having problems with the lines, wait for a while, then try the call again.

If the lines are activated and your service provider is experiencing no problems, but the wan LED is still blinking, you may have a configuration problem. Continue to the next section.

Configuration problems

If you are sure your Pipeline is properly installed, your lines are activated, and your service provider is not experiencing any problems, but the wan LED is still blinking, you may have a configuration problem.

- 1 Start your communications program and press Ctrl-L to refresh the screen.
The Configure profile appears in the Edit window:
 - 2 Check to see if you saved your Configure profile.
If an asterisk (*) appears next to Save, you have made changes to the Configure profile but did not save them. Continue to step 3.
If an asterisk does not appear next to Save, continue to step 4.
 - 3 Press Ctrl-N until the cursor moves to Save, then press Enter.
Your Configure profile is saved to the Pipeline. Try accessing the network again.
If you still have problems, continue to the next step.
 - 4 At the Configure profile, press Ctrl-D to have the Pipeline manually dial the remote site, then look at the 10-100 and 20-100 status windows to see the status of your ISDN or SW56 line:
See the *Reference Guide* chapter on status menus for more information about the messages you can see in these windows.
- If an X appears in the Link field of the 10-100 status window instead of a P, M, or D, your ISDN line is not activated or you have entered an incorrect switch type.

- If an asterisk (*) appears in the B1 or B2 field of the 10-100 status window and the remote site's name appears in the 20-100 Sessions status window, your Pipeline is connected to the remote site. Skip to step 6.
- If an asterisk (*) appears in either the B1 or B2 fields of the 10-100 status window but then disappears, any of the following configuration settings may be incorrect:
 - Rem Name: You may have entered the wrong name for the remote host.
 - Rem Addr: You may have entered the wrong IP address for the remote host.
 - Send Auth: You may have selected the wrong authentication protocol.
 - Send PW: You may have entered the password incorrectly.
 - My Name: The name you assigned to your Pipeline does not match the name expected by the remote host.
 - My Addr: The IP address you entered for your Pipeline is incorrect.
 - Check the parameters you specified in the Configure profile against those you recorded in the Configuration tables. If they match, you may need to verify the parameters with the network administrator.

Continue to step 5.

- If a D appears in either the B1 or B2 fields of the 10-100 status window, you may have entered the wrong phone number for the remote site or the wrong SPID for your ISDN line configuration. Continue to step 5.

- 5** Check the Configure profile to make sure the configuration information is entered accurately.

If you entered the information incorrectly, enter the correct information in the appropriate field of the Configure profile. Be sure to save the Configure profile.

If the information is entered correctly, make sure the information you specified is accurate:

- Contact your network administrator to confirm addresses, names, and the remote phone number.
- Contact the service provider who installed your ISDN line to confirm your SPID or SPIDs.

Once you have confirmed that all the information is entered correctly and you have saved the Configure profile, try accessing the network again. If you still have trouble, continue to step 6.

- 6 If you are routing, check to make sure you have configured your computer's IP address accurately.

Refer to your computer's manual for instructions on configuring your computer's IP address.

If you still cannot access the remote network, contact the network administrator or the Internet Service Provider you are trying to access. If this also fails, contact the Ascend Customer Service at the sites listed at the front of this guide.

Traps for BRI linkUp and linkDown

Two BRI SNMP traps, which trigger alarm events or error events, are supported on the Pipeline 130. They are linkUp and linkDown. A trap is sent by the Pipeline 130 to indicate one of the following events has occurred:

- LinkUp indicates a BRI line has been physically connected to a BRI port while the Pipeline 130 is running, or a BRI line has been initialized during the boot/cold start process.

LinkDown indicates a BRI line has been physically removed from a BRI port while the Pipeline 130 is running.

Upgrading system software

This appendix includes the following topics:

What you need to upgrade system software	C-1
Displaying the software load name	C-2
The upgrade procedure	C-3



Warning: Do not “upgrade” to an older version of software. If you use an older version of software with a new Pipeline, the unit will not function and you will need to return it to Ascend for replacement.

What you need to upgrade system software

Ascend system software is continually being enhanced to support new features and improve performance. The Pipeline is designed so that you can upgrade the system software and take advantage of these new features without returning the unit to the factory.

To upgrade the system software you need the following:

- The new system software. Contact the Ascend Technical Assistance Center for upgraded software, as described at the front of this guide.
- TFTP server software, or access to it. TFTP is required when upgrading to a fat or extended load, but can be used to upgrade to any size binary. You will need to know the host name or IP address of the host running the TFTP

server, and you will need to create a tftpboot directory to hold the binary while executing the upgrade.

- Alternatively, you can use a serial connection between a PC and the Pipeline. Use a serial connection to upload a standard-size binary. You cannot upgrade to a fat or extended load with a serial connection.

Note: The HyperTerm and Terminal programs that ship with Microsoft Windows do not reliably restore saved settings. If you are using a Macintosh communications program, Macbinary must be turned off.

Displaying the software load name

Ascend software releases are distributed in software *loads*, which are binary files that you copy to a local device and download to your Pipeline unit. Software loads vary according to functionality and target platform. The name of the software load is displayed in the Sys Options status window and in fatal error messages. The load name is an important aid to troubleshooting error conditions.

Pipeline models are abbreviated p50, p75, and p13 for the Pipeline 50, 75, and 130, respectively.

Note: For the Pipeline 85, use the same binary posted for a version 2 Pipeline 75 (which was b2.p75 at the time of this publication, but may change over time—the README file on the FTP server will guide you to the correct binary). The Pipeline 75 and the Pipeline 85 are functionally identical, except for the 4-port hub on the Pipeline 85, which does not require special software binaries.

If the software includes certain options, the name of the binary file indicates network interfaces and optional functions. These abbreviations are some that are used (see the README file on the FTP server for the latest list):

Network Interfaces

```
t    T1
e    E1
b    ISDN BRI
52   Switched 56 2 wire
54   Switched 56 4 wire
l    DSL
```

Features

```
i   IP only (OSPF - no IPX, ARA)
p   IPX only
x   X.25
a   Appletalk routing
1   Old hardware (e.g., b1.p50)
2   New hardware (e.g., b2.p75)
```

Examples

```
t.p22   Pipeline 220 T1
b2.p75  new Pipeline 50 and 75, and Pipeline 85
```

Note: When downloading the newest version of software from the Ascend FTP site (<ftp.ascend.com/pub/Software-Releases>), determine which file to download by referring to the README file associated with each sub-directory.

On your Pipeline, the current load appears in the Sys Options status window. Tab to the Sys Options window and use the down arrow to see the software load. For example:

```
00-100 Sys Option
>Access Router      ^
Load: b2.p75
Switched Installed  v
```

Also see “Pipeline checks compatibility of downloaded files” on page C-11.

The upgrade procedure

Upgrading system software is a three- or four-part process, depending on the Security profile that is currently activated. The steps required include the following:

- 1 If necessary, activate a Security profile that allows for field upgrade.
- 2 Back up your configured profiles to your computer’s hard disk.
- 3 Download the system software to the Pipeline.
- 4 Restore your Pipeline configuration.

Instructions for completing these tasks are described in this appendix. Before you go any further, check to see which version of the system software is currently installed on your Pipeline and which Security profile is activated.

To see which software version is currently running on the Pipeline, look in the Sys Option status window. Refer to the *Reference Guide* for information on using the status windows.

Activating a Security Profile

If the Security profile that is currently activated has Field Service disabled, you need to activate a Security profile with Field Service enabled to upgrade. To activate the Security profile that has Field Service enabled:

- 1 Press Ctrl-D to open the DO menu, and then press P (or select P=Password).

```
Main Edit Menu
DO
>0=ESC
P=Password
```

- 2 In the list of Security profiles, select the Security profile you want to enable. By default Field Service is enabled in the Full Access profile.

```
Main Edit Menu
Security Profile
00-301 Default
00-302
00-301 Full Access
```

The Pipeline then prompts for that profile's password.

- 3 Type the password you assigned to the profile and press Enter to accept it.

```
00-300 Security
Enter Password:
[ ]

Press > to accept
```

- 4 If you enter the right password, a message states that the password was accepted and the Pipeline is using the new security level.


```
Message #119
Password accepted.
Using new security level.
```

- 5 If the password you enter is incorrect, you are prompted again to enter the password.

This section explains how to upgrade your system software. It contains the following sections:

- Guidelines for upgrading system software
- Before you begin
- Upgrading system software with a standard load
- Upgrading system software with a fat load
- Recovering from a failed fat load upgrade
- Upgrading using the serial console
- System messages

Guidelines for upgrading system software

The following table lists the different formats for Ascend system software. How you upgrade your unit depends on the version of software you are upgrading to.

Table C-1. Format of binary loads (size comparisons)

Format of load	Size
Standard (thin)	Less than 448 Kb.
Fat	Compressed size larger than 448 Kb
Extended	Compressed size larger than 448 Kb

These restrictions apply when upgrading to the various loads:

- You must use TFTP to upload a fat or extended load.

Note: To use Trivial File Transfer Protocol (TFTP) you need a TFTP server on your computer (host) or accessible over the Ethernet. You can obtain a TFTP server from software download sites on the Internet.

- If you are upgrading your software using TFTP, you must use the `fsave` command immediately after executing the `tload` command. Failure to do so may cause your Ascend unit to lose its configuration.
- Before you can upgrade to a fat or extended load, you must first upgrade to a version of software that understands the new format, then upgrade to the fat or extended load. You can upgrade directly to a thin load (which is fat-load aware) or an extended-aware load from any version of software.

Before you begin



Caution: Uploading system software overwrites all existing profiles. Save your current Pipeline configuration before you begin. After upgrading the system software, restore the configuration. Since the saved configuration is readable text, you can manually reenter the settings, if necessary. For more information, see how to save a configuration in your Pipeline documentation.

Before upgrading your system software:

- 1 Obtain the appropriate load file, either by downloading it from the FTP server or by contacting Ascend technical support.



Caution: Be sure your unit can handle the binary; for example, an older Pipeline (with a switch on the back) cannot use a binary for a version 2 Pipeline (such as b2.p75). If you “upgrade” to a version of software not supported by your unit, the unit will no longer function and you will need to return it to Ascend for repair.

- 2 Save the current configuration.

Note: For security reasons, passwords are not included in the saved configuration text file. When you restore the configuration, the default (factory-set) passwords are reinstated. See the section on Security profiles in your documentation for more information.

- 3 If necessary, activate a Security profile that allows for field upgrade.
If you are not sure how, see the section on Security profiles in your documentation.
- 4 If you are using TFTP, be sure you have loaded the correct binaries into the / tftpboot directory on the TFTP server.

Upgrading system software with a standard load

You can upgrade system software with a standard load using either the serial console or by using TFTP over the Ethernet.

Upgrading using the serial console

- 1 From the VT100 interface, access the diagnostics monitor by typing these characters in rapid succession:
Press Ctrl-D to invoke the DO menu and select D=Diagnostics.
- 2 Enter `fsave` to save your current configuration to flash memory.
- 3 Enter `quit` to exit the Diagnostic interface.
- 4 Type the following four-key sequence in rapid succession (press each key in the sequence shown, one after the other, as quickly as possible):
Esc [Esc -
(Press the escape key, the left bracket key, the escape key, and the minus key, in that order, in rapid succession.) The following string of Xmodem control characters appear:
`CKCKCKCK`
If you do not see these characters, you probably did not press the four-key sequence quickly enough. Try again—most people use both hands and keep one finger on the escape key.
- 5 Use the Xmodem file transfer protocol to send the system file to the Pipeline.
- 6 Your communications program begins sending the file to your Ascend unit. This normally takes anywhere from 5 to 15 minutes. The time displayed on the screen does not represent real time. Do not worry if your communication program displays several “bad batch” messages. This is normal.
- 7 When the upgrade process completes, the Pipeline resets. When the self-test completes, the unit’s initial menu appears in the Edit window with all parameters set to default values.

Upgrading system software

The upgrade procedure

- 8 From the VT100 interface, access the diagnostics monitor by typing these characters in rapid succession:
Press Ctrl-D to invoke the DO menu and select D=Diagnostics.
- 9 Type `nvrampclear` to clear any differences in NVRAM memory before and after the upgrade. After the Ascend unit clears NVRAM memory, it automatically resets.
- 10 The unit resets a second time to load the configuration from flash memory.

This completes the upgrade.

Note: You can also restore your configuration from the text file saved on your hard disk. If you are not sure how to restore a configuration, see the section on restoring a configuration in the documentation.

Upgrading standard load using TFTP

- 1 Obtain the correct binary from `ftp.ascend.com/pub/Software-Releases/Pipeline`. Place the binary in a TFTP boot directory accessible via the Ethernet. Be sure the TFTP server is running. Be sure you know the IP address or host name of the server.
- 2 From the Pipeline VT100 interface, press Ctrl-D to invoke the DO menu and select D=Diagnostics.
- 3 At the `>` prompt, type:
`tload hostname filename`
where *hostname* is the name or IP address of your TFTP server (which is your computer or a server on your LAN that has a TFTP server program running), and *filename* is the name of the binary that you placed in your TFTP server's boot directory.
For example:
`tload hummer b2.p75`
or
`tload 192.168.100.2 b2.p75`
loads `b2.p75` into the Pipeline from a host named *hummer*, or loads `b2.p75` into the Pipeline from a host with an IP address of 192.168.100.2.
- 4 Enter the following command to save your configuration to flash memory:
`fsave`

- 5 Enter the following command to clear any differences in NVRAM memory before and after the upgrade.

```
nvrampclear
```

After executing this command, the Pipeline will be inaccessible while it clears NVRAM and resets. Please wait for the unit to reset before attempting to use it.

This completes the upgrade.

Upgrading system software to a fat or extended load

To upgrade your system to a fat or extended load, when your unit is currently using a standard load, you must first upgrade your system to a version that understands the new format. First upgrade to a thin load, then to a fat load.

- 1 Obtain the correct binary from <ftp.ascend.com/pub/Software-Releases/Pipeline>. Place the binary in a TFTP boot directory accessible via the Ethernet. Be sure the TFTP server is running. Be sure you know the IP address or host name of the server.
- 2 From the Pipeline VT100 interface, press Ctrl-D to invoke the DO menu and select D=Diagnostics.
- 3 At the > prompt, type:

```
tload hostname filename
```

where *hostname* is the name or IP address of your TFTP server (which is your computer or a server on your LAN that has a TFTP server program running), and *filename* is the name of the binary that you placed in your TFTP server's boot directory.

For example:

```
tload hummer b2.p75
```

or

```
tload 192.168.100.2 b2.p75
```

loads b2.p75 into the Pipeline from a host named *hummer*, or loads b2.p75 into the Pipeline from a host with an IP address of 192.168.100.2.

- 4 Enter the following command to save your configuration to flash memory:

```
fsave
```
- 5 Enter the following command to clear any differences in NVRAM memory before and after the upgrade.

```
nvramclear
```

After executing this command, the Pipeline will be inaccessible while it clears NVRAM and resets. Please wait for the unit to reset before attempting to use it.

- 6 Repeat the procedure, this time uploading the fat or extended load. Be sure your system is backed up before you begin so you can revert to a saved configuration, if necessary.

After a successful upgrade, one of the following messages appears.

- If the load is thin:

```
UART initialized
thin load: inflate
.....
...
starting system...
```

- If the load is fat:

```
UART initialized
fat load: inflate
.....
....
starting system...
```

- If the load is extended:

```
UART initialized
extended load: inflate
.....
....
starting system...
```

This completes the update load if you have no errors. If the upgrade is not successful, refer to “Recovering from a failed upgrade” next.

Recovering from a failed upgrade

If a load has an “incompatible format” message, you must first download a thin or extended-aware load that can understand the new format.

If a load has a CRC error, the following message appears:

```
UART initialized
fat load: bad CRC!!
forcing serial download at 57600 bps
please download a "thin" system...
```

Immediately after this message appears, the serial console speed is switched to 57600 bps, and the Pipeline initiates an Xmodem serial download.

To recover from this error and load the new system, you must load a thin system that is fat load aware, or an extended-aware system:

- 1 Invoke your Xmodem software to load the thin load through the console port.
- 2 After you have finished loading the prerequisite load, reboot the unit.
- 3 Download the new load using the tloadcode command.

When you download a fat load, messages similar to the following appear on the diagnostics monitor screen:

```
> tload 192.168.100.2 b2.p75
saving config to flash
.....
loading code from 192.168.100.2:69
file b2.p75..
fat load part 1:
.....
.....
fat load part 2:
.....
```

Note the “fat load part *x*:” messages. They notify you when the first and second halves of the fat load are being loaded.

Pipeline checks compatibility of downloaded files

The Pipeline compares the software to be downloaded to the currently loaded software when performing either a serial or TFTP upgrade. If the platform or network interface does not match, the Pipeline aborts the download and displays information about why the abort occurred. (The Pipeline will bypass this check if you use the TFTP command with the -f flag.)

This feature protects you from unknowingly downloading software that is incompatible with your Pipeline.

This check is initiated by the currently-loaded software. If your Pipeline is using a version of software with this feature and you attempt to load an older version of software that does not have this feature, the download will be aborted because the older software has no platform identifiers that the currently-loaded software uses to validate compatibility. In this case, you'll need to use TFTP with the -f flag, or the diagnostics command `dnldCode -f`, to have the Pipeline download the older software without performing the compatibility check.

In the following example, a user attempts to use TFTP to download a Pipeline 50 software load (b.p50) to a newer Pipeline 75 running b2.p75:

- 1 From the VT100 interface, user accesses the diagnostics monitor.
- 2 User enters the following command:

```
tload tftpserver b.p50
```
- 3 The Pipeline 75 displays the following information to the screen:

```
saving config to flash
.....
loading code from tftpserver.ascend.com
file /tftpboot/b.p50...
thin load:
This load appears to be for another platform.
This load appears not to support your network
interface
Download aborted. Use 'tloadcode -f' to force.
```

The Pipeline has compared the downloading file, b.p50 to its currently-loaded file, b2.p75. These informational messages indicate that the user attempted to load an incompatible platform and an incompatible network interface.

In the following example, a user attempts to use TFTP to download an old version of software (without this feature) to a Pipeline 75 that uses this feature:

- 1 From the VT100 interface, user accesses the diagnostics monitor.
- 2 User enters the following command:

```
tload tftpserver b.p75
```
- 3 The Pipeline 75 displays the following information to the screen:

```
saving config to flash
.....
```



```
loading code from tftpserver.ascend.com
file /tftpboot/b.p75...
thin load:
This load has no platform identifier. Proceed with
caution.
Download aborted. Use 'tloadcode -f' to force.
```

In the previous example, the user decides that he or she requires the older version and forces the download. The following messages are displayed:

- 1 User enters the following command

```
tloadcode -f tftpserver b.p75
```

- 2 The Pipeline 75 displays the following messages:

```
Download forced by user...
.....
.....
.....
```

Glossary

Authentication—A method of identifying a caller before accepting a call. The Pipeline supports token card authentication, as well as standard password, and encrypted password authentication. (Encryption is a method of encoding and decoding data.)

Bandwidth—The amount of information that can flow through a line, measured in bits per second

Bridging—One method the Pipeline can use to move data between your network and a remote network. Bridging makes remote networks look like one large network.

Channelized versus nailed—A connection can use multiple channels of available bandwidth, as in ISDN, which provides two B channels, one or two of which can be used for the same call, or you can have a permanently connected, fixed amount of bandwidth in a nailed connection.

Clearing a call—Hanging up the call gracefully. A call usually involves a number of switches. Clearing a call shuts down all the connections end to end.

Clock speed—The pace at which the Pipeline keeps synchronized with the network. The Pipeline usually takes its clock source from the network, but can generate its own clocking signal when two units are connected together.

Compression—A method of reducing the size of data to increase performance. Some algorithms maximize speed, some maximize data compression. The compress software must be present on both ends of a connection to be used.

Dynamic Bandwidth Allocation (DBA)—A proprietary method (developed by Ascend) to add or subtract B channels as needed to make the most efficient use of connection resources.

Broadcast packets—Those sent to all users on a network, even if they are for only one user. When the Pipeline is defined as a bridge, they can cause the unit to dial out.

Dialing out versus initiating a session—Anytime the Pipeline initiates a session with a remote network it dials out, but you don't have to dial or connect to Dial-Up networking, as all the dialing is done automatically. If you want to dial manually, use the DO Dial command. (See the "DO Command Reference" in the Reference Guide.)

Ethernet-to-ISDN routing—The Pipeline is an Ethernet-to-ISDN router. When you connect a Pipeline to a computer, you set up a network that uses Ethernet to carry the local network traffic. When data needs to reach a destination that is not on your local network, the data is forwarded to the Pipeline to be routed to the remote network. Before the Pipeline routes the data to the remote network, it removes the Ethernet information and repackages the data so that it can be transported over an ISDN signal through the public switched telephone network. When data comes into the Pipeline from a remote network, it extracts the data from the ISDN signal, adds Ethernet information, and places the data on your local Ethernet network.

Filter—Means to deliberately allow or disallow certain packets into the network.

Frame Relay—A service provided by the telephone company to transport data, where the line is always connected (nailed). Once the connection is established, it remains connected until either end physically disconnects the line or loses power.

IP—Internet Protocol, an addressing standard used in TCP/IP networks.

IPX—Internetwork Packet Exchange, and is used in Novell networks.

LCD interface—A term used to refer to the menu-driven Pipeline software. Originally, the menus were viewable in a palm-top Liquid Crystal Display (LCD) device. It is now referred to as the VT-100 interface because you use a VT-100 terminal emulation window to view the menus.

Packet—Refers to a block of data that has a definite order of information. Each packet contains a “packet header” that includes in it the sender’s and recipient’s address, plus the data payload and other information. Surrounding a packet is a frame, which includes information about the transport protocol.

Profile—A menu (including submenus) that defines a link or system.

Q.931 en-bloc dialing—A function included in the ISDN User-Network Interface Layer 3 Specification for Call Control, which has to do with the messages that are sent over the D channel to set up and disconnect calls.

Remote device or remote end—Refers to another network. The Pipeline dials up to or receives calls from a device at the remote end. For telecommuters, the remote end is the corporate LAN.

Routing—A method of moving data between your local network and a remote network. A router requires on-board software that enables it to deliver packets to a precise network address. Routing has many advantages over bridging, the most important being that it provides better performance.

Serial WAN port—The terminal connector on the back of some Pipeline 130 models. When wired for V.35 serial communication, the port supports a high-speed data connection to your computer from the wide area network.

Tearing down a call—See clearing a call.

Trunk groups—Lines that enable the routing of calls between switches. In the case of IDSL, the DSL service you are attached to over your ISDN line is not part of the public-switched telephone network (PSTN), but there is a route to the PSTN over a designated route, which is defined at the central office as a trunk group. The central office administrator can give you the trunk group number you need to use to direct out-going voice calls from the IDSL equipment to the PSTN.

User Datagram Protocol (UDP)—Part of the TCP/IP protocol. It was designed to provide a way for a packet to get to a particular application, rather than to a network or a host on a network. UDP uses the IP address and an additional address, called a port number. The port number for the APP Server utility is 7001.

Glossary

VT-100 terminal emulation

When the Pipeline issues a UDP unicast packet to the APP server, it sends a request to an application on a particular host, since it knows the IP address of the host, and the port number of the application. If the host doesn't have a permanent IP address, then the Pipeline broadcasts a request to all hosts on the local network. When the APP server responds, it uses the IP address of the Pipeline and the same port number, which ensures that the response goes to exactly the right process on the Pipeline.

VT-100 terminal emulation—See LCD interface.

Wide Area Network (WAN)—All remote networks not attached to the local network that you reach by connecting to a telecommunications service. The Internet as well as a remote corporate network can be referred to as the wide area network.

Index

Numerics

2nd Adrs parameter 2-15

A

ACE security 3-10, 7-18

Activation parameter 1-32, 1-36

Active parameter 1-9, 1-28

Add Pers parameter 1-8, 1-18, 1-19, 1-21

address pools 3-11

addresses

- assigning IP 3-10

- connecting bridge table to physical 5-2

- Dial Brdcast address 5-3

- netmask notation of 2-4

- routing between two IP 3-4

- spoofing local IP 6-16

- subnet 2-4

administration

- commands for performing tasks 8-10

- commands/security levels of 8-3

- features in the VT100 interface 8-1

- from a Telnet session 8-3

Adv Dialout Routes parameter 2-14

advertised routes 2-29

- poison down routes 2-15

- prevented for down routes 2-29

- redundant routes 2-14

Alt Dial#n parameter 1-9

alternate dial numbers 1-10

Alternate Mark Inversion (AMI)

- encoding limitations 1-34

AnsOrig parameter 1-5, 1-13, 1-24, 7-15

Answer profile

- bandwidth settings 1-9

- configuring for bridging connection 5-4

- how calls are answered 1-5

- ID Auth parameter in 7-14

- IPX Options submenu, configuring 4-5

- preventing a connection 1-8

- setting PPP parameters in 2-7, 2-8, 4-11

- setting up a basic profile 1-7

APP Host parameter 7-14, 7-21, A-2, A-3

APP Port parameter 7-14, 7-21

APP Server configuration

- Axent SecureNet, with A-3

- banner A-3

- DOS installation A-8

- linking to the utility A-2

- Macintosh installation A-13

- UNIX installation A-6

- Windows (all versions) installation A-10

APP Server parameter 7-14, 7-21, A-2

AppleTalk bridged connection 5-10

AppleTalk call filter, functions of 6-26

AppleTalk data filter, functions of 6-12

AppleTalk Echo Protocol (AEP) 6-12

Ascend Password Protocol (APP) A-1

Ascend Tunnel Management Protocol (ATMP)
2-37

Index

B

Ascend-Home-Agent-Password attribute 2-38
ATMP tunnels
 described 2-37
 example of how to set up 2-38
 RADIUS authenticates mobile nodes 2-37
attenuation, See Buildout parameter 1-34
Auth profile 7-14, A-2
authenticating on caller's number 1-11
authentication
 CACHE-TOKEN 7-20
 CACHE-TOKEN-CHAP 1-22
 Challenge Handshake Authentication Protocol (CHAP) 1-2
 function described 1-2
 how to assign 1-15
 Microsoft CHAP (MS-CHAP) 1-2
 PAP and CHAP 7-12
 PAP/CHAP 7-12
 PAP-TOKEN 7-19
 PAP-TOKEN-CHAP 7-20
 Password Authentication Protocol (PAP) 1-2
AUTOEXEC.NCF file 4-12
Aux Send PW parameter 7-20
average line utilization (ALU)
 calculations described 1-17

B

B channel waits before making a call 1-12
B channels used per connection 1-16
backing up, configuration 8-11
backup connection specified 1-35
BackUp parameter 1-11
BACP parameter 1-22, 1-23
bandwidth
 how to manage 1-16
 settings in the Answer profile 1-9
Bandwidth Allocation Control Protocol (BACP) 1-22
 described 1-3

 how to configure 1-22
Base Ch Count parameter 1-18, 1-20
Bill # parameter 1-13, 1-14
blinking WAN LED, troubleshooting B-10
Block calls after parameter 1-12
Blocked duration parameter 1-12
blocking connections 1-4
BOOTP
 client described 3-9
 DHCP enabled at the same time 3-10
 relay described 3-9
 server described 3-9, 3-10
BOOTP Relay profile 3-9
Bootstrap Protocol (BOOTP) 3-9
box-based routing 2-10
BRI interface, troubleshooting B-8
BRI linkUp and linkDown B-16
Bridge Adrs profile 5-7
 configuring for bridging connection 5-8
Bridge parameter 1-8, 1-9, 1-10
bridge table
 static table entries 5-8
bridge tables 5-8
 connecting to physical address 5-2
 creating/maintaining 5-7
bridged connections
 configuring 5-9
 how calls are initiated 1-4, 5-2
 planning 5-9
bridging
 globally enabling 5-2
 IPX client, to 5-13
 IPX servers, between 5-14
 parameters for 5-7
 planning connection for 5-9
 transparent 5-7
 troubleshooting problems with B-12
 used with routing 5-16
Bridging parameter 5-6
broadcast address described 2-6

broadcast addresses, from Dial Brdcast 5-3

broadcast packets

- initiate bridged connections 1-4

- setting them to dial out or not 1-11

Buildout parameter 1-34

C

CACHE-TOKEN-CHAP authentication 1-22

call filter

- AppleTalk 6-26

- described 6-4

- IP 6-26

- NetWare 6-21

Call Filter parameter 1-11, 6-4

Call Type parameter 1-13, 1-24, 1-28, 1-29

Callback parameter 1-13, 7-15, 7-16

Called # parameter 1-9, 1-11

Calling # parameter 1-9, 7-15

calling back the caller to authenticate 1-13

calling line ID authentication, see Id Auth parameter 1-8

Calling-line ID 7-13

calls

- authenticating incoming 7-11

- authenticating using PAP and CHAP 7-12

- clearing calls 8-15

- manually placing/clearing 8-10

- preventing initiation of 6-21

- problems with T1 line B-11

- See also

 - Answer profile

 - Connection profile

 - connections

- troubleshooting B-2

calls delayed

- see Preempt parameter 1-12

Channel Service Unit (CSU) enabled 1-33

channels used for a connection 1-16

CHAP described 7-12

clearing idle MP+ calls 1-19

CLID (Calling Line ID)

- configuring 7-13

CLID Fail Busy parameter 7-14

Client Gateway parameter 2-20

Clock Source parameter 1-33, 1-34

clock speed of the serial WAN port 1-35

COM port, setting Term Rate to same as B-5

Com port, setting Term Rate to same as 8-5

commands

- accessing administration 8-3

- displaying terminal server 8-17

- for administrative tasks 8-10

- security/manual tasks of DO 8-3

- Sys Reset 8-15

- terminal server 8-2

Compare parameter 6-9

compression methods supported 1-6

configuration

- APP Server utility A-2

- bridged connections 5-9

- Filter profiles 6-6

- IPX SAP filters 4-7

- NetWare clients 4-10

- NetWare LANs 4-25

- of DNS addresses 3-7

- restoring 8-14

- system 8-4

Connection # parameter 4-7, 4-18, 5-7

connection cannot be reached

- see Secondary and Backup parameters 1-12

connection charges, bill to number 1-14

Connection profile

- deactivating 1-10

- defining individual connections 1-9

- first one created from Configure menu 1-2

- setting encapsulation 1-14

- setting Session options 1-11

- setting Telco options 1-13

- Static Rte profiles and 2-17

connection security 7-11

Index

D

connections
 bridge IP 5-16
 bridging AppleTalk 5-10
 bridging IPX client 5-13
 bridging IPX server 5-14
 configuring IP address for 2-33
 configuring IPX routing 4-24
 configuring RIP for 2-23
 configuring RIP for incoming WAN 2-22
 manually placing 8-11
 network-to-network 2-33
 processes following established 8-21
 routing IP 2-3
 static IPX routes 4-18
 See also bridged connections
console interface, type specified 8-5
Console parameter 8-5
Contact parameter 8-4
cost management, call filters used for 6-4
Customer Premises Equipment (CPE), discussed for Frame Relay 1-30

D

Data Circuit-Terminating Equipment (DCE)
 1-36
data compression
 example settings 1-16
 MS-Stac 1-7
 Stac 1-6
 VJ Comp 1-7
Data Filter parameter 1-11, 6-3
data filters
 described 6-2
 used for security 7-16
Data Link Connection Identifier (DLCI) 1-25
Data Svc parameter 1-13, 1-27, 1-28, 1-29
DBA Monitor parameter 1-19, 1-20
DCE N39n parameters 1-28
deactivating a Connection profile 1-10

Default (security) profile 7-4
default gateway, Rem Adrs parameter and 2-19
default route
 configuring 2-19
 defining per user 2-20
 setting RIP to ignore 2-22
 usage 2-19
default security
 changing 7-1
 level recommendations 7-4
Default security profile 7-4
Dest parameter 2-15, 2-16, 2-18
DHCP
 BOOTP enabled at the same time 3-10
 client 3-25
 Server 3-10, 3-23
 how to set up 3-14
 Spoofing
 how to set up 3-14
 menu 3-11
 response 3-10
DHCP Spoofing profile 3-11
Dial # parameter 1-9, 7-15
Dial Brdcast address 5-3
Dial Brdcast parameter 1-5, 1-10, 1-11, 5-7
Dial Query parameter 1-5, 4-8
Dial Query, functions of 4-8
dial-in NetWare clients 4-5
dial-in Windows 95 clients 4-4
dialing
 manually placing/clearing 8-10
 problems with manual 8-11, B-2
dial-out packets displayed 6-28, 6-29
disable routing of incoming packets 3-32
disconnect cause code
 on authentication failure 7-14
 value set 7-14
DLCI parameter 1-27, 1-31
DNS (Domain Name System)
 configuring for 3-7

DNS host address table 3-17
 DNS list attempt 3-7
 DNS profile 3-2, 3-15
 Dnstab edit command 8-18
 Dnstab entry command 8-18
 Dnstab show command 8-18
 DO commands
 accessing 8-3
 availability B-2
 for security/manual tasks 8-3
 using 8-10
 DO commands described 8-3, 8-10
 Domain Name Server (DNS), assigned 3-15
 DS0 channel assignments 1-33
 Dst Adrs parameter 6-10
 Dst Mask parameter 6-10
 Dst Network Adrs parameter 4-22
 Dst Node Adrs parameter 4-22
 Dst Port # parameter 6-10
 Dst Port Cmp parameter 6-10
 Dst Socket # parameter 4-23
 Dst Socket Cmp parameter 4-23
 DTE N39n parameters 1-28
 dual IP 3-4
 Dyn Alg parameter 1-8, 1-17, 1-21
 Dynamic Bandwidth Allocation (DBA) 1-16, 1-17
 Dynamic Host Configuration Protocol (DHCP) 3-10, 3-25
 dynamic IP routing 3-4
 dynamic routes
 described 2-16, 2-30
 enabling 2-21

E

Edit Security parameter 7-8
 Edit System parameter 7-9

Encaps options 1-14
 Encaps parameter 1-9, 1-15, 1-20, 1-31
 encapsulation
 setting in the Connection profile 1-10
 Encoding parameter 1-33
 Enet Adrs parameter 5-7
 Ether Options submenu 2-15, 3-2
 Ethernet interface
 assigning IP address to 3-3
 configuring IP routing 3-1
 configuring IPX routing 4-12
 turning on bridging 5-6
 Exp Callback parameter 7-16
 Expect Callback
 90-second wait enabled 7-16

F

Facilities Data Link (FDL) 1-33
 FDL parameter 1-33
 Field Service parameter 7-9
 Filter Persistence parameter 1-11, 6-36, 6-37
 Filter profile
 components of 6-6
 defining/applying 6-16
 predefined 6-21
 filter vs. firewall persistence 6-36
 filters
 AppleTalk data filter 6-12
 call 6-4
 example generic 6-12
 example IP filter 6-16, 6-19
 example IPX filter 4-22
 example IPX RIP 6-23
 NetWare call 6-21
 numbers for 6-4
 persistence described 1-12
 firewalls
 assigned to a Connection profile 6-35
 configured for port routing 3-29

Index

G

First DSO parameter 1-33
Force 56 parameter 1-7, 1-13, 1-14
Forward parameter 4-22, 6-9
FR Prof parameter 1-27, 1-31
FR Type parameter 1-28, 1-29
fractional T1 1-13
Frame Relay
 configuring 1-28
 forcing the link up at all times 1-29
 gateway connection example 1-30
 Gateway mode 1-26
 obtaining the DLCI number 1-31
 option installed 1-25
 planning for 1-25
 shown in status window 1-25
 status reports 1-30
Frame Relay profile 1-25
Framing Mode parameter 1-32
FT1 Caller parameter 1-13, 1-24
Full Access profile 7-3
Full Access profile, activating 8-3

G

Gateway parameter 2-17, 2-18
gateway, Rem Adrs parameter as default 2-19
Generic filter
 conditions for 6-9
 described 6-8
Generic Routing Encapsulation (GRE) 2-37
group numbers
 assigning 1-3
 how to assign 1-4
Group parameter 1-13, 1-24, 1-36

H

Handle IPX parameter 4-8

Handle IPX Type20 parameter 4-8
Hangup command 8-18
hardware configuration, troubleshooting B-3
Help command 8-18
Home Agent ATMP end point 2-37
home agent configured in router mode 2-38
Hop Count parameter 4-6, 4-18
host route subnet address requirements 2-6
hosts
 software requirements for IP 2-30
 software requirements for IPX 4-10
 using PPP dial-in software 2-2
 with their own IP network 2-2

I

ICMP (Internet Control Message Protocol)
 Redirect packets, function of 2-21
ICMP redirects 7-7
ICMP Redirects parameter 2-14, 2-21, 7-7
Id Auth parameter 1-5, 1-7, 1-8, 7-14
idle link not used
 see Preempt parameter 1-12
Idle parameter 1-11, 1-22, 6-5
Idle Pct parameter 1-8, 1-19, 1-21
idle timer
 function described 1-12, 6-4
 preventing resetting of 6-21
 reset by RIP updates 2-23
IF Adrs parameter 2-11, 2-12
Ignore Def Rt parameter 2-15, 2-22
incoming calls
 assigning dynamic address to 2-2
 authenticating 7-11
 authenticating on phone number 1-11
 IP routing requirements 2-2
input filter
 conditions described 6-6
 of IP call filter 6-26

-
- SAP filters 4-7
 - interface-based routing
 - described 2-10, 2-13
 - Internet Group Membership Protocol (IGMP) 2-13
 - Inverse Address Resolution Protocol (InARP)
 - response data described 1-32
 - supported for Frame Relay 1-31
 - inverse multiplexing 1-16
 - IP address
 - assigned automatically 3-10
 - assigning to Ethernet interface 3-3
 - preventing spoofing in a filter 7-16
 - subnet mask notation 2-4
 - IP Adrs parameter 2-15
 - IP bridged connection 5-16
 - IP filter
 - conditions for 6-10
 - described 6-8
 - IP Options submenu 2-11, 2-20, 3-16
 - IP routing
 - and ICMP Redirects 2-21
 - and Ping command 3-6
 - and RIP version 2 2-9
 - assigning two interface addresses 3-4
 - authenticating 2-2
 - configurations 2-30
 - configured with bridging 2-2
 - configured with IPX routing 2-2
 - configuring a static host route 2-31
 - configuring subnet 3-3
 - Default route 2-19
 - host requirements listed by platform 2-30
 - overview of 2-1
 - parameters enabling 3-3
 - planning configuration for 2-30
 - second destination specified 2-30
 - sharing dynamic (dual IP) 3-4
 - static routes 2-16
 - UDP checksums 3-8
 - IP routing table
 - built at start up 2-8
 - choosing which routes go in the table 2-24
 - description of all fields 2-27
 - example configurations 2-30
 - how host routes are added 2-2
 - how network routes are added 2-3
 - management 2-14
 - preventing large tables 2-22
 - temporary routes 2-30
 - updating local router's 3-8
 - usage described 2-8
 - viewing the routing table 2-25
 - IP subnet addresses 2-4
 - Iproute add command 2-8, 8-18
 - Iproute command 8-18
 - Iproute delete command 8-18
 - Iproute show command 2-25, 8-18
 - IPX
 - client bridging 5-13
 - filter for RIP packets 6-23
 - filters 4-6
 - Ping command 4-14
 - RIP broadcasts, controlling 4-17
 - routes dropped 4-7
 - server bridging 5-14
 - IPX Enet# parameter 4-5, 4-13
 - IPX Frame parameter 4-13
 - IPX network 4-6
 - configuration with servers and clients 4-25
 - with only servers 4-29
 - IPX Options submenu 4-8
 - IPX Pool# parameter 4-5, 4-15
 - IPX RIP parameter 4-8, 4-17
 - IPX RIP table 4-16
 - IPX route and SAP packets dropped 4-7
 - IPX Routes profile 4-4, 4-6, 4-18
 - IPX routing
 - and authentication of callers 4-11
 - client considerations 4-10
 - configuring 4-24
 - configuring IPX SAP on a WAN link 4-20
 - connecting a dial-in user 4-4
-

Index

L

-
- defining a network for dial-in clients 4-15
 - dynamic addresses for dialin clients 4-4
 - enabling system-wide 4-11
 - extensions for WAN links 4-3
 - filtering SAP packets 4-7
 - learning the Ethernet IPX number 4-14
 - local NetWare server issues 4-12
 - NetWare client software 4-10
 - NetWare server table 4-2
 - NetWare server table displayed 4-16
 - RIP default route 4-3
 - Routing table displayed 4-16
 - SAP filters 4-2
 - using IPX RIP for dynamic routes 4-3
 - watchdog spoofing 4-9
 - IPX Routing parameter 4-13
 - IPX SAP broadcasts, controlling 4-20
 - IPX SAP Filter parameter 1-11, 4-11
 - IPX SAP Filter profile 4-21
 - IPX SAP filters 4-21, 4-23
 - IPX SAP filters profile 4-7
 - IPX SAP parameter 4-8, 4-20
 - IPX SAP Proxy Net#n parameter 4-14
 - IPX SAP table 4-16
 - IPX server bridge 5-14
 - IPX Type 20 packets 4-8
 - IPXping command 4-14, 8-18
 - ISDN BRI lines
 - circuit-quality problems with B-10
 - troubleshooting B-8
 - troubleshooting problems with B-9
- ### L
- LAN Adrs parameter 1-6, 2-11, 2-12, 2-15, 2-16, 2-17, 2-33, 2-34, 2-36
 - Last DSO parameter 1-33
 - learning bridge 5-7
 - LEDs
 - troubleshooting blinking WAN B-10, B-13
 - Length parameter 6-9
 - Link Comp parameter 1-8, 1-15, 1-16, 1-21
 - link encapsulation supported
 - Bandwidth Allocation Control Protocol (BACP) 1-3
 - Frame Relay RFC 1490 1-3
 - Multichannel Point-to-Point Protocol (MPP) 1-3
 - Multilink PPP (MP) 1-2
 - Multilink Protocol Plus (MP+) 1-3
 - PPP 1-2
 - link management protocol 1-30
 - Link Mgmt parameter 1-27, 1-28, 1-30
 - Link Quality Management (LQM) reporting
 - periods set 1-16
 - linkDown B-16
 - links, problems with quality of B-12
 - linkUp B-16
 - LinkUp parameter 1-28, 1-29
 - List attempt 3-7
 - Local command 8-17, 8-18
 - local DNS table 3-17
 - configuration 3-18
 - creating 3-19
 - deleting 3-21
 - editing 3-20
 - local management information, configuring for 8-4
 - Location parameter 8-4
 - Log Facility parameter 8-6
 - Log Host parameter 8-6
 - log messages, working with 8-9
 - Log Port parameter 8-6
 - Log profile 8-6
 - logical link configurations described 1-26
 - LOGIN.EXE 4-10
 - Loopback parameter 1-33
 - LQM parameters 1-8, 1-15, 1-16, 1-21
-

M

MAC (Ethernet) addresses 3-10
 Macintosh clients of NetWare servers 4-10
 manually dialing a connection 1-4
 manually dialing, problems with 8-11, B-2
 Mask parameter 6-9
 Max Ch Count parameter 1-8, 1-19, 1-20, 1-24
 Maximum Receive Unit (MRU) packet size 1-15
 messages, working with status/log 8-9
 Metric parameter 2-18
 Min Ch Count parameter 1-8, 1-18, 1-20
 Module Enabled parameter 1-36
 monitoring DBA 1-19
 More parameter 6-9
 MP connections described 1-16
 MP+
 configuring a profile with 1-20
 connections described 1-16
 cost considerations 1-19
 tasks to set up a connection 1-17
 MRU parameter 1-8, 1-15, 1-16, 1-20, 1-29
 Multicast Forwarding parameter 2-13
 multicast forwarding, described 2-13
 Multicast submenu 2-13
 multi-channel links described 1-16
 Multichannel Point-to-Point Protocol (MPP), described 1-3
 Multilink PPP (MP), described 1-2
 Multilink Protocol Plus (MP+), described 1-3
 multiple-address NAT 3-25

N

N391 parameter 1-28
 nailed connections

 channel limits discussed 1-3
 described 1-3, 1-13
 determining which end adds bandwidth 1-14
 shown in routing table 2-29
 Nailed Grp parameter 1-27, 1-28, 1-29
 Nailed T1 Group parameter 1-32
 Nailed T1 profile 1-32
 Nailed/MPP connections 1-23, 1-24
 Name Binding Protocol (NBP) 6-12
 Name parameter 1-15, 1-28
 naming the Remote end of the connection 1-10
 NAT 3-23
 DHCP requests 3-26
 for Frame Relay 3-28
 multiple-address translation 3-25
 profile 3-30
 single-address translation 3-29
 translation table size 3-24
 NAT profile 3-26
 Net Adrs parameter 5-7
 netmask notation 2-4
 netmask values of subaddresses 2-5
 NetWare
 see IPX routing
 NetWare call filter, functions of 6-21
 Netware t/o parameter 4-8
 network address translation (NAT) 3-22
 network base address described 2-6
 network bits in subnet addresses 2-5
 Network number, used to reach an IPX network 4-6
 Network parameter 4-6, 4-18
 Node parameter 4-6, 4-18
 Normal call clearing disconnect cause code 7-14
 number of channels used for a connection 1-16
 number to dial out, where to enter 1-10
 numbered interfaces 2-11

Index

O

O

Offset parameter 6-9

Operations parameter 7-8

origin of connection settings 1-13

output filter

- conditions described 6-6

- in NetWare Call 6-22

- of IP call filter 6-26

- SAP filters 4-7

P

Packet Burst 4-10

packets

- defining filter types for 6-8

- dial-out, displayed 6-29

- disable routing of 3-32

- forwarding/blocking 6-2

- identifying outbound SAP 6-22

PAP-TOKEN authentication

- for outbound calls 7-19

PAP-TOKEN-CHAP authentication

- for outbound calls 7-20

Passwd parameter 7-9

Password Authentication Protocol (PAP) 7-12

passwords

- default full access 7-5

- for establishing bridging 5-3

- hidden in Security profiles 7-8

- how verified 7-13

- recommended initial changes 7-2

- SNMP 7-5

- Telnet 7-2

Peer parameter 4-5, 4-8

permanent switched connection 1-13

physical addresses, keeping track of 5-7

Ping command 8-18

Plug and Play 3-10

- how to set up 3-14

Point-to-Point Protocol (PPP), described 1-2

poison dialout routes when a link is down 2-15

port numbers of common ports 3-24

port routing 3-29

- configuration 3-29

Power-On Self Test (POST) 8-16

PPP dial-in software used by host 2-2

PPP encapsulation 1-14

PPP negotiation 3-23

PPP-encapsulated call authentication 7-12

Preempt parameter 1-11

preferred servers, NetWare configurations for
4-10

prefixes removed from called number 1-11

preventing unwanted connections 1-4

primary connection lost 1-35

private addresses vs. official addresses 3-23

Private parameter 2-19

privileges in Security profiles 7-8

Profile Reqd parameter 1-5, 1-7, 1-8, 7-7

propagating RIP and SAP packets 4-14

Protocol parameter 6-10

protocols

- AARP (AppleTalk Address Resolution) 6-12

- AEP (AppleTalk Echo Protocol) 6-12

- APP (Ascend Password Protocol) A-1

- BOOTP 5-1

- IPX 4-1, 4-2, 4-3

- link-level bridging 5-1

- PPP IPXC 4-1

- SAP (Service Advertising Protocol) 4-3

- TCP/IP 6-10, 6-11

Q

Quit command 8-18

R

R/W Comm Enable parameter 7-6
 R/W Comm parameter 7-6
 Read Comm parameter 7-6
 rebooting device 8-2
 receiving an incoming call 1-5
 Recv Auth parameter 1-5, 1-8, 1-9, 2-7, 7-12
 Recv PW parameter 1-15
 Red Alarm state, troubleshooting B-10
 redial attempts controlled
 Block Calls After parameter 1-12
 redundant routes advertised 2-14
 Registered Ports 3-32
 Remote command 8-18
 remote interface address 2-12
 remote management
 setting higher terminal rate for 8-5
 Remote management via Telnet 8-3
 Remote Mgmt parameter 8-5
 reserved IP addresses 3-10
 resetting the unit 8-15
 Restore Cfg 8-12, 8-15
 restoring saved configurations 8-14
 RIP (Routing Information Protocol) 2-21
 configuring for a connection 2-23
 configuring for incoming WAN connections 2-22
 configuring on local Ethernet 2-22
 default route for IPX 4-3
 filter for IPX RIP packets 6-23
 for dynamic IP routing 2-9
 IPX RIP 4-3
 recommendations for use 2-21
 static routes and 2-17
 RIP and SAP, related to dial-in clients 4-6
 RIP parameter 2-7, 2-15, 2-21, 2-22, 2-23
 RIP Policy parameter 2-14
 Rip Preference parameter 2-25

RIP Summary parameter 2-14
 RIP v1 as historic 2-21
 RIP version 2 support 2-9
 RIP2 Use Multicast parameter 2-15
 Route IP parameter 1-8, 1-9, 1-10, 2-7
 Route IPX parameter 1-8, 1-9, 1-10, 4-5
 route metrics discussed 2-24
 route preferences
 listed by route type 2-24
 router
 updating on the backbone 3-8
 routing
 between NetWare LANs 4-1
 enabling dynamic 2-21
 stop advertising down routes 2-15
 table limitations for IPX servers 4-7
 using IP 2-1
 routing connections
 how calls are initiated 1-4
 routing, used with bridging 5-16

S

SAFEWORD 3-10, 7-18
 SAP filters 4-2
 SAP packets
 dropped 4-7
 identifying outbound 6-22
 prevented from initiating a call 1-12
 SAP Service Type 4-19
 Save Cfg 8-12
 Sec History parameter 1-8, 1-17, 1-18, 1-19, 1-21
 Secondary parameter 1-12
 Secure Access Firewall software 6-34
 Secure Access Manager (SAM) 6-35
 security
 activating 7-4, 8-3
 default enabled after reset 7-5

Index

S

- default level 7-4, 7-8
- defining new Security profiles 7-10
- full access level 7-3
- ICMP redirects off 7-7
- password authentication features 7-12
- passwords in Security profiles 7-3
- privileges 7-8
- privileges in Full Access profile 7-9
- recommended measures 7-1
- security cards, using 7-18
- Security menu 7-3
- Security profiles
 - activating Field Service C-4
- security profiles 7-8, 7-10
 - activating 7-4
 - upgrading, used when C-4
- Send Auth parameter 1-14, 1-15, 7-19
- Send PW parameter 1-15, 7-19
- Serial WAN profile 1-35
- Server Name parameter 4-7, 4-18
- Server Type parameter 4-6, 4-7, 4-18
- servers
 - NetWare configurations for preferred 4-10
- Service Type parameter 4-19
- Session options 1-11
- Set all command 8-18
- Set ARP clear command 8-18
- Set FR command 8-18
- Set password command 7-22, 8-19
- Set sessid command 8-19
- Set term command 8-19
- Show ARP command 8-19
- Show DHCP address 8-19
- Show DHCP command 8-19
- Show DHCP lease 8-19
- Show dnstab command 8-19
- Show dnstab entry command 8-19
- Show FR DLCI command 8-19
- Show FR lmi command 8-19
- Show FR stats command 8-19
- Show ICMP command 8-19
- Show IF stats command 8-19
- Show if totals command 8-19
- Show igmp clients command 8-19
- Show igmp groups command 8-19
- Show igmp stats command 8-19
- Show ip address command 8-19
- Show ip routes command 8-20
- Show ip stats command 8-20
- Show isdn command 8-20
- Show netw networks commands 8-20
- Show netw pings command 8-20
- Show netw servers command 8-20
- Show netw stats command 8-20
- Show netware command 4-16
- Show revision command 8-20
- Show sessid command 8-20
- Show tcp connection command 8-20
- Show tcp stats command 8-20
- Show udp listen command 8-20
- Show udp stats command 8-20
- Show uptime command 8-20
- single-address NAT, configuring 3-29
- slash notation of subnet masks 2-5
- SNEP (Serialization Number Exchange Protocol) 6-24
- SNMP community strings 7-5
- SNMP management, described 8-3
- SNMP Options submenu 7-6
- Socket parameter 4-6, 4-18
- software load name C-2
- Split Code.User parameter 1-21, 1-22
- spoofing, address 6-16
- SPX spoofing 4-9
- SPX watchdog 4-9
- Src Adrs parameter 6-10

Src Mask parameter 6-10
 Src Network Adrs parameter 4-22
 Src Node Adrs parameter 4-22
 Src Port # parameter 6-10
 Src Port Cmp parameter 6-10
 Src Socket # parameter 4-22
 Src Socket Cmp parameter 4-22
 Stac data compression supported 1-6
 static bridge table entries 5-8
 Static Mappings profile 3-27
 Static Preference parameter 2-25
 static route described 2-16
 Static Rtes profile 2-17, 2-18
 Connection profile and 2-17
 station names, for establishing bridging 5-3
 Station parameter 1-9
 status information, access to 8-2
 status messages, working with 8-9
 status windows 8-9
 Sub Pers parameter 1-8, 1-18, 1-19, 1-21
 subnet addresses
 described 2-4
 table of values 2-4
 switched call type described 1-13
 Sys Config menu 8-4
 Sys Config profile 1-14
 Sys Diag menu 8-12
 Sys Diag menu, described 8-10
 Sys Reset 8-12, 8-16
 Sys Reset command, described 8-15
 syslog messages
 format 8-7
 from the firewall 8-7
 Syslog parameter 8-6
 Syslog, configuring 8-2, 8-5
 system device 8-2
 system events, maintaining permanent log of
 8-5

system security
 activating 7-4
 for Telnet 7-5
 system-based routing 2-10

T

T1 interface, troubleshooting problems with
 B-9
 T1 line
 circuit-quality problems with B-10
 how to configure 1-32
 planning 1-32
 troubleshooting access problems B-11
 T1.617D link management settings 1-30
 T391 and T392 parameters 1-28
 Target Util parameter 1-8, 1-18, 1-21
 Tcp command 8-20
 TCP Estab parameter 6-10
 TCP ports 3-23
 Telco Options 1-13
 Telnet command 8-20
 Telnet password required 8-21
 Temporary parameter 2-29
 Term Rate
 setting to higher rate for remote management
 8-5
 setting to the same speed as COM port B-5
 setting to the same speed as Com port 8-5
 Term Rate parameter 8-5
 Term Serv 8-12
 Term Serv menu 8-17
 terminal server
 accessing command-line 8-17
 commands for 8-2
 terminal server commands 8-17
 Test command 8-20
 TFTP command C-11
 Tick Count parameter 4-7, 4-18

Index

U

- timers in filters 6-4
- token card authentication
 - and Split Code.User parameter 1-22
 - by multiple users 1-22
 - configured with APP Server A-1
 - RADIUS authentication server used with 1-22
- Traceroute command 8-20
- transparent bridging 5-7
- troubleshooting
 - configuration B-2
- troubleshooting problems
 - for bridge/router B-12
 - for hardware configuration B-3
 - general types of B-2
 - ISDN BRI interface B-8
 - T1 and ISDN BRI circuit-quality B-10
 - T1 and ISDN BRI interface B-9
- trunk number removed from called number 1-11
- tsave -a command 8-13
- tunnels, configuring ATMP 2-37

U

- UDP checksums 3-7, 3-8
- UDP Port parameter 2-38
- UDP ports 3-23
- UDP/IP sessions described 2-37
- UNIX clients for NetWare servers 4-10
- upgrading on-board software C-3
- User Busy disconnect cause code 7-14

V

- V.35 serial WAN port, configuring 1-35
- Value parameter 6-9
- version of current software C-2

- virtual private networks 2-37
- VJ Comp data compression supported 1-7
- VJ Comp parameter 1-8, 1-15, 1-16, 1-21
- VT100 control terminal
 - hardware configuration with B-3

W

- wait before retrying to redial
 - Blocked Duration parameter 1-12
- WAN Alias parameter 2-11, 2-12, 2-15
- WAN connections
 - configuring RIP for 2-22
 - Filter profile connected to 6-16
- WAN LED, troubleshooting blinking B-10
- watchdog spoofing, described 4-9
- wdDialout diagnostic command 6-29
- weighting algorithm used to set up DBA 1-17
- Well Known Ports 3-32
- Windows NT required for MS-CHAP 1-15

Z

- zero address of a subnet mask 2-7