

Pipeline 220 User's Guide

*Ascend Communications, Inc.
Part Number: 7820-0323-001
For Software Version 5.1Ap5*

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

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

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

Ascend Customer Service

You can request 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
- Whether you are routing or bridging with your Ascend product
- 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
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>

Important safety instructions

The following safety instructions apply to the Pipeline 220:

- 1 Read and follow all warning notices and instructions marked on the product or included in the manual.
- 2 The maximum recommended ambient temperature for Pipeline 220 models is 104° Fahrenheit (40° Celsius). Make sure to allow sufficient air circulation or space between units when the Pipeline 220 is installed in a closed or multi-unit rack assembly, because the operating ambient temperature of the rack environment might be greater than room ambient.
- 3 Openings in the cabinet are provided for ventilation. To ensure reliable operation of the product and to protect it from overheating, these openings must not be blocked or covered.
- 4 Installation of the Pipeline 220 in a rack without sufficient air flow can be unsafe.
- 5 If installed in a rack, the rack should safely support the combined weight of all equipment it supports. The Pipeline 220 weighs 7.25 lbs (3.3 kg).
- 6 The connections and equipment that supply power to the Pipeline 220 should be capable of operating safely with the maximum power requirements of the Pipeline 220. In the event of a power overload, the supply circuits and supply wiring should not become hazardous. The input rating of the Pipeline 220 is printed on its nameplate.
- 7 Models with AC power inputs are intended for use with a three-wire grounding type plug—a plug which has a grounding pin. This is a safety feature. Equipment grounding is vital to ensure safe operation. Do not defeat the purpose of the grounding type plug by modifying the plug or using an adapter.
- 8 Before installation, use an outlet tester or a voltmeter to check the AC receptacle for the presence of earth ground. If the receptacle is not properly grounded, the installation must not continue until a qualified electrician has corrected the problem.
- 9 If a three-wire grounding type power source is not available, consult a qualified electrician to determine another method of grounding the equipment.

-
- 10** Install only in restricted-access areas in accordance with Articles 110-16, 110-17, and 110-18 of the National Electrical Code, ANSI/NFPA 70.
 - 11** Do not allow anything to rest on the power cord, and do not locate the product where persons will walk on the power cord.
 - 12** Do not attempt to service this product yourself. Opening or removing covers can expose you to dangerous high voltage points or other risks. Refer all servicing to qualified service personnel.
 - 13** General purpose cables are provided with this product. Special cables, which might be required by the regulatory inspection authority for the installation site, are the responsibility of the customer.
 - 14** When installed in the final configuration, the product must comply with the applicable Safety Standards and regulatory requirements of the country in which it is installed. If necessary, consult with the appropriate regulatory agencies and inspection authorities to ensure compliance.
 - 15** A rare phenomenon can create a voltage potential between the earth grounds of two or more buildings. If products installed in separate buildings are *interconnected*, the voltage potential might cause a hazardous condition. Consult a qualified electrical consultant to determine whether or not this phenomenon exists and, if necessary, implement corrective action before interconnecting the products.

In addition, if the equipment is to be used with telecommunications circuits, take the following precautions:

- Never install wiring during a lightning storm.
- Never install jacks in wet locations unless the jack is specifically designed for wet locations.
- Never touch uninsulated wires or terminals unless the telephone line has been disconnected at the network interface.
- Use caution when installing or modifying lines.
- Avoid using equipment connected to lines (other than a cordless telephone) during an electrical storm. There is a remote risk of electric shock from lightning.
- Do not use a telephone or other equipment connected to telephone lines to report a gas leak in the vicinity of the leak.

About This Guide

This guide explains how to install, configure, and administer the Pipeline 220. It also explains how to navigate the Java-based Ascend Configurator, which runs on any Windows 95 or Windows NT 4.0 computer.

How to use this guide

This guide contains the following chapters:

- Chapter 1, “Introduction” provides an overview of Pipeline 220 applications, Pipeline 220 configuration, and Pipeline 220 management features.
- Chapter 2, “Installing the Pipeline 220,” explains how to install the Pipeline 220 hardware.
- Chapter 3, “Using the graphical interfaces,” explains how to install the Ascend Configurator on any Windows 95 or Windows NT computer and navigate its interface.
- Chapter 4, “Configuring the Pipeline 220 for WAN Access,” explains how to configure the Pipeline 220’s WAN interface.
- Chapter 5, “Configuring Frame Relay”, describes Frame Relay connections and how to configure them.
- Chapter 6, “Configuring IP Routing,” describes IP routing and how to configure it.
- Chapter 7, “IP Address Management”, describes BOOTP Relay, DHCP, and Network Address Translation (NAT).
- Chapter 8, “Configuring OSPF Routing,” describes OSPF routing and how to configure it.
- Chapter 9, “Setting Up IP Multicast Forwarding,” describes Multicast forwarding and how to configure it.
- Chapter 10, “Configuring IPX Routing,” describes IPX routing and how to configure it. Included is a detailed discussion of NetWare RIP and SAP packets.
- Chapter 11, “Configuring AppleTalk Routing,” describes AppleTalk routing and how to configure it. Included are detailed discussions of AppleTalk zones and network numbering.
- Chapter 12, “Configuring Packet Bridging,” describes bridging, how to configure it, and how to manage the bridge table.
- Chapter 13, “Defining Static Filters,” explains static packet filtering and how to configure it on the Pipeline 220.
- Chapter 14, “Setting Up Virtual Private Networking,” explains the network tunneling concept, Ascend Tunnel Management Protocol (ATMP), and how to configure ATMP.
- Chapter 15, “SNMP administrative support,” describes the SNMP protocol, and how to configure it, including configuration of Ascend traps.
- Chapter 16, “VT100 Interface System Administration,” explains the Pipeline 220 VT100 interface and how you can use it to monitor status and manage the Pipeline 220.


What you should know

This guide is for the person who configures and maintains the Pipeline 220. To configure the Pipeline 220, you need to understand the following:

- Internet or telecommuting concepts
- Wide area network (WAN) concepts
- Local area network (LAN) concepts

Documentation conventions

This section explains all the special characters and typographical conventions 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.
<i>Italics</i>	Represent variable information. Do not enter the words themselves in the command. Enter the information they represent. In ordinary text, italics are used for titles of publications, for some terms that would otherwise be in quotation marks, and to show emphasis.
[]	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.
	Separates command choices that are mutually exclusive.
>	Points to the name of an item you select from a menu. This symbol appears between the name of a menu and the name of the item you should select from the menu. (The <i>menu</i> does not necessarily appear at the top of the screen. For example, you might open it by clicking a button.)
Key1-Key2	Represents a combination keystroke. To enter a combination keystroke, press the first key and hold it down while you press one or more other keys. Release all the keys at the same time. (For example, Ctrl-H means hold down the Control key and press the H key.)
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.

Convention

Meaning



Warning:

Warns that a failure to take appropriate safety precautions could result in physical injury.

Manual set

The Pipeline 220 Documentation Set consists of the following manuals:

- Pipeline 220 User's Guide (this manual)
- Pipeline 220 VT100 Interface Guide
Explains how to use the VT100 user interface to configure the Pipeline 220.
- Pipeline 220 VT100 Interface Reference Guide.
Describes each command in the VT100 user interface.

Related RFCs

RFCs are available on the Web at <http://ds.internic.net>.

Information about PPP connections

For information about PPP connections and authentication, you might want to download one or more of the following:

- RFC 2153: *PPP Vendor Extensions*
- RFC 1994: *PPP Challenge Handshake Authentication Protocol (CHAP)*
- RFC 1990: *The PPP Multilink Protocol (MP)*
- RFC 1989: *PPP Link Quality Monitoring*
- RFC 1974: *PPP Stac LZS Compression Protocol*
- RFC 1962: *The PPP Compression Control Protocol (CCP)*
- RFC 1877: *PPP Internet Protocol Control Protocol Extensions for Name Server Addresses*
- RFC 1662: *PPP in HDLC-like Framing*
- RFC 1661: *The Point-to-Point Protocol (PPP)*
- RFC 1638: *PPP Bridging Control Protocol (BCP)*
- RFC 1332: *The PPP Internet Protocol Control Protocol (IPCP)*

Information about IP routers

RFCs that describe the operation of IP routers include:

- RFC 2030: *Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI*
- RFC 2002: *IP Mobility Support*
- RFC 1812: *Requirements for IP Version 4 Routers*

- RFC 1787: *Routing in a Multi-provider Internet*
- RFC 1519: *Classless Inter-Domain Routing (CIDR): An Address Assignment and Aggregation Strategy*
- RFC 1433: *Directed ARP*
- RFC 1393: *Traceroute Using an IP Option*
- RFC 1256: *ICMP Router Discovery Messages*

Information about OSPF routing

For information about OSPF routing, see:

- RFC 1850: *OSPF Version 2 Management Information Base*
- RFC 1587: *The OSPF NSSA Option*
- RFC 1586: *Guidelines for Running OSPF Over Frame Relay Networks*
- RFC 1583: *OSPF Version 2*
- RFC 1246: *Experience with the OSPF protocol*
- RFC 1245: *OSPF protocol analysis*

Information about multicast

For information about multicast, see:

- RFC 1949: *Scalable Multicast Key Distribution*
- RFC 1584: *Multicast Extensions to OSPF*
- RFC 1458: *Requirements for Multicast Protocols*

Information about packet filtering

RFCs that describe firewalls and packet filters include:

- RFC 1858: *Security Considerations for IP Fragment Filtering*
- RFC 1579: *Firewall-Friendly FTP*

Information about general network security

RFCs pertinent to network security include:

- RFC 1704: *On Internet Authentication*
- RFC 1636: *Report of IAB Workshop on Security in the Internet Architecture*
- RFC 1281: *Guidelines for the Secure Operation of the Internet*
- RFC 1244: *Site Security Handbook*

ITU-T recommendations

ITU-T recommendations (formerly CCITT) are available commercially. You can order them at <http://www.itu.ch/publications/>.

Related publications

This guide and documentation set do not provide a detailed explanation of products, architectures, or standards developed by other companies or organizations.

Here are some related publications that you might find useful:

- William Flanagan. *The guide to T1 Networking*.
- Uyless Black. *Data Link Protocols*
- W. Richard Stevens. *TCP/IP Illustrated*
- William R. Cheswick and Steven M. Bellovin. *Firewalls and Internet Security*

Contents

Ascend Customer Service	iii
Important safety instructions	iv

About This Guide..... vii

How to use this guide	vii
What you should know	viii
Documentation conventions	viii
Manual set	ix
Related RFCs	ix
Information about PPP connections	ix
Information about IP routers	ix
Information about OSPF routing	x
Information about multicast	x
Information about packet filtering	x
Information about general network security	x
ITU-T recommendations	x
Related publications	xi

Chapter 1 Introduction..... 1-1

Using the Pipeline 220 for private and public access	1-1
Who uses the Pipeline 220?	1-1
What are some common applications of the Pipeline 220?	1-1
Dual LAN access	1-2
WWW access for all Internet users	1-2
Virtual Private Networking (VPN)	1-3
Internet Gateway	1-4
Overview of Pipeline 220 configuration	1-4
Creating a network diagram	1-4
Configuring lines, slots, and ports for WAN access	1-5
Configuring WAN connections and security	1-5
Concentrating Frame Relay connections	1-5
Configuring routing and bridging across the WAN	1-6
Protocol-independent packet bridging	1-6
IPX routing (NetWare 3.11 or newer)	1-6
IP routing	1-6
Configuring Internet services	1-6
Multicast	1-6
OSPF routing	1-6
Virtual Private Networking (VPN)	1-7
Overview of management features	1-7
Using the Ascend Configurator	1-7
Using the terminal server command line	1-7

Using status windows to track WAN or Ethernet activity	1-8
Managing the Pipeline 220 by means of SNMP	1-8
Flash RAM and software updates	1-8
Where to go next	1-9

Chapter 2 Installing the Pipeline 220..... 2-1

What is included with the Pipeline 220?	2-1
Additional required hardware	2-2
WAN interface	2-2
Computer with a serial port	2-2
Modem cable	2-2
Ethernet interface	2-2
Required software	2-3
Networking software	2-3
Communications software	2-3
Installation overview	2-4
Choosing a location for the Pipeline 220	2-5
Connecting the Pipeline 220 to the computer's Ethernet interface	2-5
Connecting to an Ethernet network	2-6
Connecting a computer to the Pipeline 220	
Terminal port	2-6
Connecting an IBM-compatible computer	2-6
Connecting a Macintosh	2-7
Connecting a Unix workstation	2-7
Connecting the Pipeline 220 to your leased line	2-8
Starting up the Pipeline 220	2-8
Interpreting the Pipeline 220 LEDs	2-9

Chapter 3 Using the graphical interfaces..... 3-1

Before you begin	3-1
Installing the Ascend Configurator	3-2
Using the Configurator Startup screen	3-3
Settings on the Startup screen	3-3
Using the Options screen	3-4
Settings on the General tab screen	3-4
Settings on the Network tab screen	3-5
Settings on the Syslog tab screen	3-6
Using the Ascend Configurator interface	3-7
Ascend Configurator tabs	3-9
If you are familiar with Ascend's VT100 user interface	3-9
Assigning the Pipeline 220 an IP address	3-10
Connecting to a Pipeline 220	3-13
Opening an existing configuration file	3-13
Uploading a configuration	3-13
Saving a configuration to a file	3-14
Security issues with the Configurator	3-14
Using the Pipeline 220 Quickstart	3-14

Chapter 4 Configuring the Pipeline 220 for WAN Access..... 4-1

Introduction to WAN configuration	4-1
---	-----

Configuring the T1 line	4-1
T1 parameters	4-2
T1 line framing and encoding	4-2
Amount of attenuation required	4-2
Clock source for synchronous transmission	4-2
Configuring the nailed T1 line	4-2
Entering the settings	4-2
Enabling the settings	4-4
Configuring the serial WAN port	4-4
Signals to control the serial WAN data flow	4-4
Configuring the serial WAN interface	4-4
 Chapter 5 Configuring Frame Relay	5-1
Using the Pipeline 220 as a Frame Relay concentrator	5-1
Kinds of physical network interfaces	5-2
Kinds of logical interfaces to a Frame Relay switch	5-2
Network to Network Interface (NNI)	5-2
User to Network Interface — Data Communications Equipment (UNI-DCE)	5-2
User to Network Interface — Data Terminal Equipment (UNI-DTE)	5-2
Types of Frame Relay connections	5-3
Gateway connections	5-3
Frame Relay circuits	5-3
Configuring the logical link to a Frame Relay switch	5-4
Understanding the Frame Relay parameters	5-4
Specifying a Frame Relay profile name and activating the profile	5-4
Bringing down the datalink when DLCIs are not active	5-4
Defining the nailed connection to the switch	5-4
Specifying the type of Frame Relay interface	5-4
Link management protocol	5-4
Frame Relay timers and event counts	5-5
Maximum Receive Units (MRU)	5-5
Examples of Frame Relay profile configurations	5-5
Configuring an NNI interface	5-5
Saving the settings	5-6
Configuring a UNI-DCE interface	5-7
Configuring a UNI-DTE interface	5-8
Configuring Connection profiles for Frame Relay	5-9
Understanding the Frame Relay connection parameters	5-9
Gateway connections	5-9
Frame Relay circuits	5-9
Examples of connection configurations	5-9
Configuring a Frame Relay gateway connection	5-9
Configuring a Frame Relay circuit	5-13
 Chapter 6 Configuring IP Routing.....	6-1
Introduction to IP routing and interfaces	6-1
IP addresses and subnet masks	6-1
Zero subnets	6-3
IP routes	6-4
How the Pipeline 220 uses the routing table	6-4
Static and dynamic routes	6-4

Route preferences and metrics	6-5
Pipeline 220 IP interfaces	6-5
Ethernet interfaces	6-5
WAN IP interfaces	6-6
Numbered interfaces	6-6
Configuring the local IP network setup	6-8
Understanding the IP network parameters	6-8
Primary IP address for each Ethernet interface	6-8
Second IP address for each Ethernet interface	6-8
Enabling RIP on the Ethernet interface	6-9
Ignoring the default route	6-9
Proxy ARP and inverse ARP	6-9
Telnet password	6-10
BOOTP relay	6-10
Local domain name	6-10
DNS or WINS name servers	6-10
DNS lists	6-10
Client DNS	6-10
SNTP service	6-11
Specifying SNTP server addresses	6-11
UDP checksums	6-11
Poisoning dialout routes in a redundant configuration	6-11
Examples of IP network configurations	6-12
Configuring the Pipeline 220 IP interface on a subnet	6-12
Joining a subnet	6-12
Saving the settings	6-13
Making the backbone router the default route	6-14
Configuring DNS	6-15
Configuring IP routing connections	6-17
Understanding the IP routing connection parameters	6-17
Enabling IP routing for WAN connections	6-17
Enabling IP routing for a WAN interface	6-17
Configuring the remote IP address	6-17
WAN alias	6-18
Specifying a local IP interface address	6-18
Assigning metrics and preferences	6-18
Private routes	6-18
Configuring RIP policy on the WAN interface	6-19
Checking remote host requirements	6-19
UNIX software	6-19
Windows or OS/2 software	6-19
Macintosh software	6-19
Software configuration	6-19
Examples of IP routing connections	6-20
Configuring a router-to-router connection	6-20
Configuring a router-to-router connection on a subnet	6-22
Configuring a numbered interface	6-24
Configuring IP routes and preferences	6-27
Understanding the static route parameters	6-27
Route names	6-27
Activating a route	6-27
Route's destination address	6-27

Route's gateway address	6-28
Virtual hops, costs, and preferences	6-28
Tagging routes learned from RIP	6-28
Type-1 or type-2 metrics for routes learned from RIP	6-28
Making a route private	6-28
Routes for Connection profile interfaces	6-29
A connected route for the Ethernet IP interface	6-29
Static route preferences	6-29
RIP and OSPF preferences	6-29
Examples of static route configurations	6-29
Configuring the default route	6-29
Defining a static route to a remote subnet	6-31
Example route preferences configuration	6-33
Configuring the Pipeline 220 for dynamic route updates	6-34
Understanding the dynamic routing parameters	6-34
RIP (Routing Information Protocol)	6-34
Ignoring the default route	6-34
RIP policy and RIP summary	6-34
Ignoring ICMP redirects	6-35
Private routes	6-35
Examples of RIP and ICMP configuration	6-35
Configuring RIP policy	6-35
Configuring RIP on the WAN link	6-36
Syslog services	6-37
Configuring the Pipeline 220 to send Syslog messages	6-37
Syslog messages	6-40

Chapter 7 IP Address Management 7-1

BOOTP Relay	7-1
Saving your settings	7-2
DHCP services	7-3
How IP addresses are assigned	7-4
Configuring DHCP services	7-4
Enabling DHCP services	7-4
Configuring IP address pools	7-6
Assigning specific addresses to particular hosts	7-6
Local DNS host address table	7-6
User-definable TCP connection retry timeout	7-8
Network Address Translation (NAT)	7-9
NAT and port routing	7-9
Configuring NAT	7-10
Configuring NAT port routing	7-11
Well-known ports	7-12

Chapter 8 Configuring OSPF Routing 8-1

Introduction to OSPF	8-1
RIP limitations solved by OSPF	8-1
Ascend implementation of OSPF	8-2
OSPF features	8-2
Security	8-3
Support for variable length subnet masks	8-3

	Interior Gateway Protocol (IGP)	8-3
	Exchange of routing information	8-4
	Designated and backup designated routers	8-4
	Configurable metrics	8-5
	Hierarchical routing (areas)	8-6
	Stub areas	8-6
	The link-state routing algorithm	8-7
	Configuring OSPF routing in the Pipeline 220	8-10
	Understanding the OSPF routing parameters	8-10
	Examples of adding the Pipeline 220 to an OSPF network	8-11
	Configuring OSPF on the Ethernet interface	8-12
	Configuring OSPF across the WAN	8-16
	Configuring a WAN link that doesn't support OSPF	8-18
Chapter 9	Setting Up IP Multicast Forwarding	9-1
	Overview	9-1
	Understanding the multicast parameters	9-2
	Enabling multicast forwarding	9-2
	Specifying the MBONE interface	9-2
	Monitoring the multicast heartbeat	9-2
	Configuring multicast forwarding on a client interface	9-3
	An implicit priority setting for dropping multicast packets	9-3
	Forwarding from an MBONE router on Ethernet	9-4
	Configuring system-wide multicast parameters	9-4
	Configuring multicasting on WAN interfaces	9-6
	Saving the settings	9-7
	Forwarding from an MBONE router on a WAN link	9-8
	Configuring the Pipeline 220 to respond to multicast clients	9-8
	Configuring the MBONE interface	9-8
	Configuring multicasting on WAN interfaces	9-9
Chapter 10	Configuring IPX Routing	10-1
	Introduction to Ascend IPX routing	10-1
	IPX Service Advertising Protocol (SAP) tables	10-1
	IPX RIP (Routing Information Protocol) tables	10-2
	Ascend extensions to standard IPX	10-2
	IPX Route profiles	10-3
	IPX SAP filters	10-3
	WAN considerations for NetWare client software	10-3
	IPX in the Answer profile	10-4
	Enabling IPX routing	10-4
	Enabling authentication	10-5
	Applying an IPX SAP Filter to the Answer profile	10-6
	Saving the settings	10-7
	Integrating the Pipeline 220 into the local IPX network	10-8
	Checking local NetWare configurations	10-8
	Configuring IPX on the Pipeline 220 Ethernet interface	10-8
	Working with the RIP and SAP tables	10-10
	Viewing the RIP and SAP tables	10-10
	Restricting RIP in a Connection profile	10-10
	Configuring static IPX routes	10-11

	Restricting SAP in a Connection profile	10-14
	Filtering SAP traffic	10-15
	Defining an IPX SAP filter	10-15
	Applying IPX SAP filters	10-17
	Example of an IPX routing connection	10-19
	Configuring the Pipeline 220 at site A	10-19
	Enabling IPX routing for site A's Ethernet interface	10-21
	Configuring a static route from site A to the remote server	10-22
	Configuring the Pipeline 220 at site B	10-24
	Enabling IPX routing for site B's Ethernet interface	10-25
	Configuring a static route at site B	10-25
Chapter 11	Configuring AppleTalk Routing	11-1
	Introduction to AppleTalk routing	11-1
	When to use AppleTalk routing	11-1
	Reducing broadcast and multicast traffic	11-1
	Providing dynamic startup information to local devices	11-2
	Understanding AppleTalk zones and network ranges	11-2
	AppleTalk zones	11-2
	Extended and non-extended AppleTalk networks	11-2
	How AppleTalk works	11-4
	Configuring AppleTalk routing	11-5
	System-level AppleTalk routing parameters	11-5
	Saving the settings	11-6
	Per-connection AppleTalk routing parameters	11-7
Chapter 12	Configuring Packet Bridging	12-1
	Introduction to Ascend bridging	12-1
	Disadvantages of bridging	12-1
	How a bridged WAN connection is initiated	12-2
	Physical addresses and the bridge table	12-2
	Broadcast addresses	12-2
	How the Pipeline establishes a bridging connection	12-3
	Enabling bridging	12-3
	Enabling bridging on the Ethernet interface	12-3
	Saving the settings	12-4
	Bridging in the Answer Profile	12-5
	Managing the bridge table	12-7
	Transparent bridging	12-7
	Static bridge table entries	12-7
	Example of a bridged connection	12-10
	Assigning a name	12-10
	Configuring a bridging connection	12-11
	Defining a static bridge-table entry	12-11
	Configuring proxy mode on the Pipeline 220	12-12
Chapter 13	Defining Static Filters	13-1
	Introduction to Ascend filters	13-1
	How filters work	13-1
	How to apply filters	13-1

Applying a filter to the Answer profile	13-2
Saving the settings	13-3
Applying a filter to a Connection profile	13-4
Applying a filter to the Ethernet interface	13-4
Overview of Filter profiles	13-6
Filtering inbound and outbound packets	13-7
Selecting filter type, activating the filter, and defining the conditions	13-7
Defining generic filter conditions	13-8
Defining IP filter conditions	13-8
Examples of filters	13-9
Example of a generic filter to handle AppleTalk broadcasts	13-9
Setting up the filter	13-10
Blocking AARP packets	13-11
Allowing non-AppleTalk traffic to pass	13-13
Allowing AppleTalk AEP packets to pass	13-14
Blocking broadcast AppleTalk packets	13-14
Blocking NBP packets	13-15
Discarding unrecognized packets	13-16
Saving the filter	13-16
Example of an IP filter to prevent address spoofing	13-16
Setting up the filter	13-17
Discarding incoming packets with the local subnet as the source address	13-18
Discarding packets with a loopback source address	13-19
Forwarding incoming packets that have nonlocal source addresses	13-19
Forwarding outbound packets that have local source addresses	13-19
Saving the filter	13-20
A sample IP filter for more complex security issues	13-20
Setting up the filter	13-20
Forwarding packets destined for the web server	13-21
Forwarding responses to TCP requests	13-21
Forwarding UDP packets	13-22
Saving the filter	13-22

Chapter 14 **Setting Up Virtual Private Networking 14-1**

Introduction to Virtual Private Networking (VPN)	14-1
Configuring ATMP tunnels	14-1
How the Pipeline 220 creates ATMP tunnels	14-2
Router and gateway mode	14-2
Configuring a home agent in router mode	14-3
Understanding the ATMP router mode parameters	14-3
Notes about routing to the mobile node	14-4
Example of configuring a home agent in router mode (IP)	14-4
Example of configuring a home agent in router mode (IPX)	14-8
Configuring a home agent in gateway mode	14-9
Understanding the ATMP gateway mode parameters	14-10
Example of configuring a home agent in gateway mode (IP)	14-10
Example of configuring a home agent in gateway mode (IPX)	14-13

Chapter 15 **SNMP administrative support..... 15-1**

Introduction	15-1
Configuring SNMP access security	15-1

How the SNMP security options work	15-1
Community strings	15-2
Address security	15-2
Entering SNMP security settings	15-2
Setting SNMP traps	15-5
Understanding the SNMP trap parameters	15-5
Entering an SNMP trap configuration	15-5
Ascend Enterprise traps	15-8
Alarm events	15-8
Port state change events	15-8
Security events	15-9
Supported MIBs	15-10
 Chapter 16	
VT100 Interface System Administration	16-1
Introduction to Pipeline 220 administration	16-1
Accessing the VT100 interface	16-2
Using the Pipeline 220 control port	16-2
Using Telnet	16-2
Using the VT100 interface	16-3
Activating a menu or status window	16-3
Opening menus and profiles	16-3
Opening edit fields	16-4
Setting enumerated parameters	16-5
Saving your changes	16-5
Special display characters and keys	16-5
About Pipeline 220 passwords	16-7
Using the Pipeline 220 status windows	16-8
Line status window	16-9
System Events	16-9
Sessions	16-10
Dyn Stat	16-11
WAN Stat	16-12
Ether Stat	16-13
Sys Options	16-13
Ether Opt	16-14
Syslog	16-15
Terminal-server command-line interface	16-16
Exiting the terminal server interface	16-17
Commands not supported on the Pipeline 220	16-17
Commands for use by terminal-server users	16-18
Set command	16-18
Show command	16-18
IProute command	16-35
DNSTab command	16-40
Menu command	16-43
Telnet command	16-43
TCP command	16-45
Ping command	16-46
IPXping command	16-47
Traceroute command	16-48
Kill command	16-49

Appendix A	Pipeline 220 Specifications.....	A-1
	General specifications	A-1
	Battery	A-1
	Power requirements	A-2
	Environmental requirements	A-2
	User interface specifications	A-3
	Ethernet interface specifications	A-3
	Serial WAN cabling specifications	A-4
	V.35 cable to WAN	A-4
	RS-449 cable to WAN	A-5
Appendix B	Upgrading System Software	B-1
	Upgrading system software	B-2
	Definitions and terms	B-2
	Guidelines for upgrading system software	B-3
	Before you begin	B-4
	Upgrading system software	B-5
	Using TFTP to upgrade to a standard load	B-6
	Using TFTP to upgrade to a fat or thin load	B-6
	Recovering from a failed fat load upgrade	B-8
	Upgrading software with an extended load	B-9
	Upgrading software from versions earlier than 4.6C to version 5.0A or above	B-10
	Using the serial port to upgrade to a standard or a thin load	B-11
	System messages	B-14
Appendix C	Warranties and FCC regulations	C-1
	Product warranty	C-1
	Warranty repair	C-1
	Out-of warranty repair	C-2
	FCC Part 15	C-2
	FCC Part 68 Notice	C-2
	IC CS-03 Notice	C-3
	Glossary.....	Glossary-1

Figures

Figure 1-1	Dual LAN access for employees in a corporation	1-2
Figure 1-2	Public access for Internet users and private access for employees.....	1-2
Figure 1-3	Network tunneling across the Internet.....	1-3
Figure 1-4	Using the <i>Pipeline 220</i> as a central-site Internet gateway	1-4
Figure 1-5	<i>Pipeline 220</i> User's Guide Roadmap.....	1-9
Figure 2-1	Back panel of the <i>Pipeline 220</i>	2-4
Figure 2-2	Connecting an adapter to a Macintosh modem cable	2-7
Figure 2-3	Front panel of the <i>Pipeline 220</i>	2-9
Figure 3-1	Ascend Configurator interface.....	3-7
Figure 4-1	Back panel of the <i>Pipeline 220</i>	4-1
Figure 5-1	The <i>Pipeline 220</i> operating as a Frame Relay concentrator	5-1
Figure 5-2	Network to Network interface (NNI) in a <i>Pipeline 220</i> unit	5-2
Figure 5-3	User to Network Interface-Data Communications Equipment (UNI-DCE). ..	5-2
Figure 5-4	User to Network Interface - Data Terminal Equipment (UNI-DTE)	5-3
Figure 5-5	NNI interface to another switch.....	5-5
Figure 5-6	UNI-DCE interface to an end-point (DTE)	5-8
Figure 5-7	UNI-DTE interface to a Frame Relay switch	5-8
Figure 5-8	Gateway connections	5-10
Figure 5-9	A Frame Relay circuit.....	5-13
Figure 6-1	A class C IP address	6-2
Figure 6-2	A 29-bit subnet mask and number of supported hosts.....	6-2
Figure 6-3	Interface-based routing example.....	6-6
Figure 6-4	Sample dual IP network.....	6-8
Figure 6-5	Creating a subnet for the <i>Pipeline 220</i>	6-12
Figure 6-6	A router-to-router IP connection	6-20
Figure 6-7	A connection between local and remote subnets	6-22
Figure 6-8	Example numbered interface	6-24
Figure 6-9	Two-hop connection that requires a static route when RIP is off.....	6-31
Figure 8-1	Autonomous system border routers	8-3
Figure 8-2	Adjacency between neighboring routers	8-4
Figure 8-3	Designated and backup designated routers.....	8-4
Figure 8-4	OSPF costs for different types of links.....	8-5
Figure 8-5	Dividing an AS into areas.....	8-6
Figure 8-6	Sample network topology	8-7
Figure 8-7	A sample OSPF setup	8-12
Figure 9-1	<i>Pipeline 220</i> forwarding multicast traffic to multicast clients.....	9-4
Figure 9-2	<i>Pipeline 220</i> as a multicast forwarder on Ethernet and WAN interfaces	9-8
Figure 10-1	A connection with NetWare servers on both sides.....	10-19
Figure 11-1	AppleTalk LAN.....	11-3
Figure 11-2	Routed connection	11-4
Figure 12-1	Negotiating a bridge connection (PPP encapsulation).....	12-3
Figure 12-2	How the <i>Pipeline 220</i> creates a bridging table	12-7
Figure 12-3	Example of a bridged connection	12-10

Figures

Figure 13-1	Filter terminology	13-6
Figure 14-1	ATMP tunnel across the Internet.....	14-2
Figure 14-2	Home agent routing to the home network	14-3
Figure 14-3	Home agent in gateway mode.....	14-10

Tables

Table 2-1	<i>Pipeline 220</i> ports	2-4
Table 2-2	<i>Pipeline 220</i> LEDs	2-9
Table 6-1	IP address classes and default subnet masks	6-2
Table 6-2	Standard subnet masks	6-3
Table 8-1	Link state databases for network topology in Figure 8-6	8-8
Table 8-2	Shortest-path tree and resulting routing table for Router-1	8-9
Table 8-3	Shortest-path tree and resulting routing table for Router-2	8-9
Table 8-4	Shortest-path tree and resulting routing table for Router-3	8-9
Table 16-1	Special purpose keys for Control Monitor displays (<i>continued</i>)	16-6
Table 16-2	16-6
Table 16-3	Session status characters	16-10
Table 16-4	Link quality values	16-11
Table 16-5	Ether Stat fields	16-13
Table 16-6	Sys Options information	16-14
Table 16-7	Ether Opt fields	16-15
Table 16-8	Sample terminal-server menu	16-43
Table A-1	<i>Pipeline 220</i> source power requirements	A-2
Table A-2	Control Monitor port cabling pinouts	A-3
Table A-3	V.35 cable pinouts	A-4
Table A-4	RS-449 cable pinouts	A-5
Table B-1	Ascend system software versions	B-4
Table B-2	Before upgrading	B-4
Table B-3	System software messages	B-14

Introduction

This chapter covers the following topics:

Using the Pipeline 220 for private and public access	1-1
Overview of Pipeline 220 configuration	1-4
Overview of management features	1-7
Where to go next	1-9

Using the Pipeline 220 for private and public access

The Pipeline 220 is a high-performance WAN router that can enable all Internet users to access your FTP site, World Wide Web site, and any other publicly available resources, while your employees have secure access to your corporate network backbone.

The Pipeline 220 delivers WAN access through either an unchannelized T1/FT1 or a V.35 interface. You cannot use both WAN interfaces simultaneously. Your software configuration activates one or the other.

Who uses the Pipeline 220?

The most common users of the Pipeline 220 are medium to large companies, major corporations, and ISPs who provide both open access and secure access through multiple Ethernet LAN segments. The unit's configuration options provide the flexibility and security you need for optimizing your installation. Management features include a comprehensive set of control and monitoring functions and easy-to-perform upgrades.

What are some common applications of the Pipeline 220?

Some common applications that use the features provided by the Pipeline 220 are dual LAN access, secure/open Internet access, network tunneling, and operation as an Internet gateway.

Introduction

Using the Pipeline 220 for private and public access

Dual LAN access

Figure 1-1 shows a typical configuration for a company that offers publicly accessible network resources to all employees and restricted access to a separate secure network

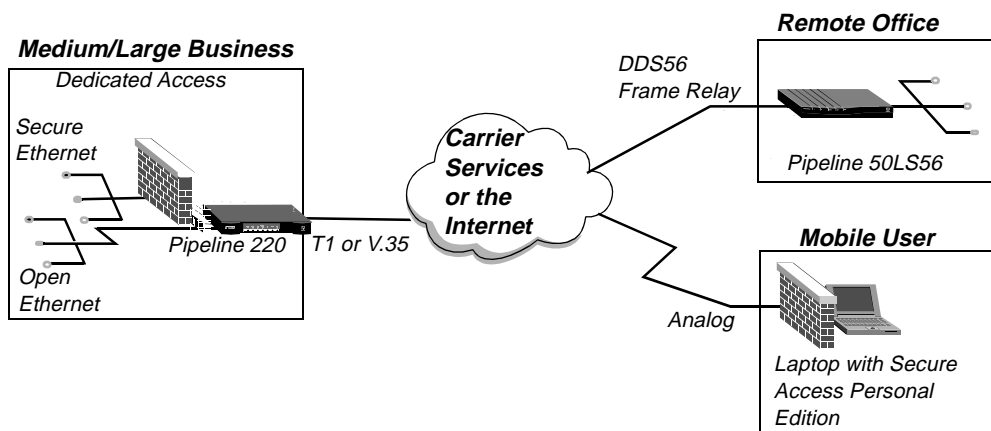


Figure 1-1. Dual LAN access for employees in a corporation

The figure shows two types of users: a remote office that connects to the corporate network through a leased Frame Relay connection, and a remote user who dials into the Internet through an ISP connection and connects to the corporate network by means of the Internet.

WWW access for all Internet users

Figure 1-2 shows a typical configuration for a company that offers a World Wide Web site and FTP site to any Internet user and a secure corporate network connection to its employees

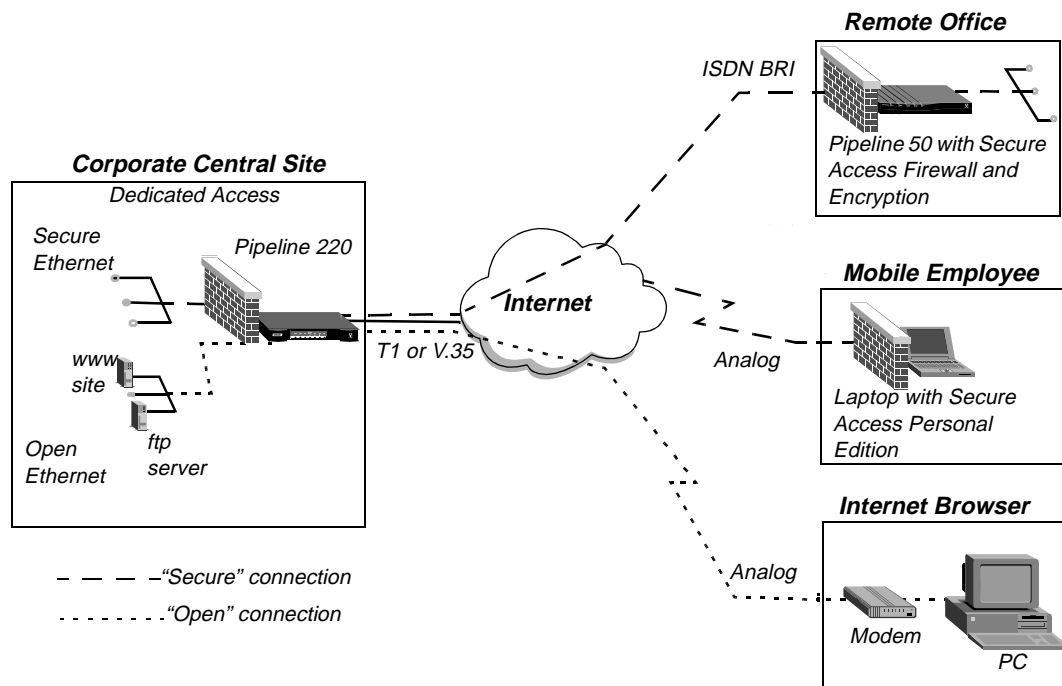


Figure 1-2. Public access for Internet users and private access for employees

You can ensure security by using built-in static filters, Connection profiles, and RADIUS, or you can use the optional Secure Access feature to ensure the highest level of security through Ascend's Secure Access Firewalls and IPSec encryption.

Virtual Private Networking (VPN)

With the optional VPN feature, you can forward non-IP traffic across the Internet. The feature supports the following protocols:

- Point To Point Tunneling Protocol (PPTP)

Note: The Pipeline 220 supports PPTP by routing or forwarding PPTP traffic as appropriate. The Pipeline 220 does not act as either a PPTP Access Concentrator (PAC) or a PPTP Network Server (PNS).

- Ascend Tunnel Management Protocol (ATMP)

Figure 1-3 shows a typical network tunneling environment.




Figure 1-3. Network tunneling across the Internet

Network tunneling adds another level of security by encrypting the data it sends across the Internet. The supported tunneling protocols can be used in combination with Secure Access.

Internet Gateway

Figure 1-4 shows a company offering switched access to corporate resources for remote offices, telecommuters, and mobile users through an Ascend MAX. The company offers its employees protected connection to the Internet through the Pipeline 220.

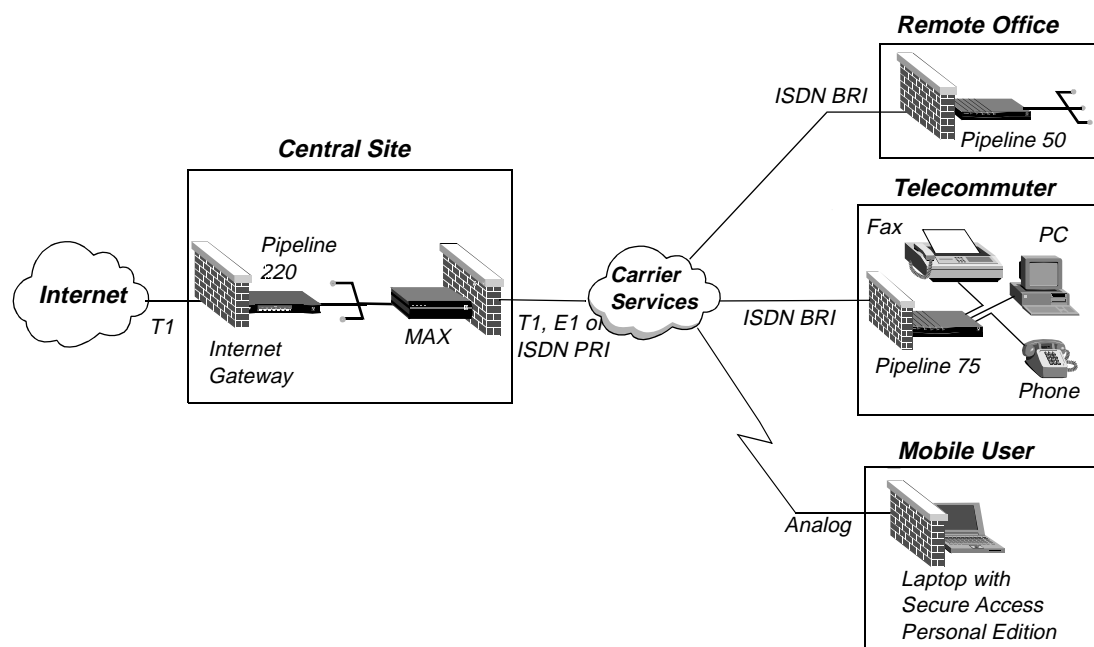


Figure 1-4. Using the Pipeline 220 as a central-site Internet gateway

Overview of Pipeline 220 configuration

This section provides an overview of how to configure the Pipeline 220. It covers the following topics:

- Configuring the lines, channels, and ports, and how calls are routed between them
- Configuring wide area network connections and security
- Configuring the Pipeline 220 as a Frame Relay or X.25 concentrator
- Configuring routing and bridging across the WAN
- Configuring Internet services, such as multicast, OSPF, and virtual private networks

Creating a network diagram

Ascend strongly recommends that, after you have read this introductory material, you diagram your network and refer to the diagram while configuring the Pipeline 220.

Creating a comprehensive network diagram helps prevent problems during installation and configuration, and can help you troubleshoot problems later.

Configuring lines, slots, and ports for WAN access

The Pipeline 220 comes with one built-in T1 or E1 line and a V.35 serial port (8 Mbps). You cannot simultaneously use both types of access. Use the Pipeline 220 Configurator to activate either the T1/E1 port or the V.35 port.

You can use either type of access for a leased high-speed connection to a Frame Relay switch or to another WAN router. Neither type requires extensive configuration. Your service provider will provide you with the small amount of information you need to configure the Pipeline 220 for WAN access. You specify most of the required information in a Frame Relay or Connection profile.

Once you have enabled the lines and ports for WAN access, you need to configure the manner in which users are routed to them, for access across the WAN, and routed from them to other destinations (such as the local network).

Configuring WAN connections and security

When the Pipeline 220 receives packets that require routing to a remote network, it forwards them across the WAN connection. Software at the both ends of the connection encapsulates each packet before sending it out over the WAN. Each type of encapsulation supports its own set of options, which can be configured on a per-connection basis to enable the Pipeline 220 to interact with a wide range of software and devices.

After a connection's link encapsulation method has been negotiated, the Pipeline 220 typically uses a password to authenticate the call. Following are some of the connection security features supported in the Pipeline 220:

Feature	Description
Authentication protocols	For PPP connections, the Pipeline 220 supports both Password Authentication Protocol (PAP) and Challenge-Handshake Authentication Protocol (CHAP). CHAP is more secure than PAP, and is preferred if both sides of the connection support it.
Terminal server security	After a dial-in user has passed the initial-connection security check, another password can be required for access to the Pipeline 220 terminal services. Within the terminal server, you can restrict which commands are accessible to users, or prevent users from executing any command other than Telnet.
Filters and firewalls	Filters and firewalls provide a packet-level security mechanism that can provide a very high level of network security.

Concentrating Frame Relay connections

The Pipeline 220 provides extensive support for Frame Relay. Using a T1 line or serial WAN port for a nailed connection to a switch, it can function as a Network-to-Network Interface (NNI) switch, a Data Communications Equipment (DCE) unit responding to users, or a Data Terminal Equipment (DTE) unit requesting services from a switch.

Configuring routing and bridging across the WAN

Routing and bridging configurations enable the Pipeline 220 to forward packets between the local network and the WAN.

Protocol-independent packet bridging

The Pipeline 220 can operate as a link-level bridge, forwarding packets from Ethernet to a WAN connection (and vice versa) on the basis of the destination hardware address in each packet. Unlike a router, a bridge does not examine packets at the network layer. It simply forwards packets to another network segment if the address does not reside on the local segment.

IPX routing (NetWare 3.11 or newer)

The Pipeline 220 can operate as an IPX router, linking remote NetWare LANs with the local NetWare LAN on Ethernet.

IP routing

IP routing is the most widespread use of the Pipeline 220, and the unit has a wide variety of configurable options. IP routing is the required basis for Internet-related services such as IP multicast support, OSPF, and cross-Internet tunneling for virtual private networks. Most sites create static IP routes to enable the Pipeline 220 to reliably connect to certain destinations or to change global metrics or preferences settings.

Configuring Internet services

All Internet services and routing methods require that the Pipeline 220 function as an IP router, so an IP routing configuration is a necessary precondition.

Multicast

The multicast backbone (MBONE) is a virtual network layered on top of the Internet to support IP multicast routing across point-to-point links. It is often used for transmitting audio and video on the Internet in real time, because multicasting is a much cheaper and faster way to communicate the same information to multiple hosts.

OSPF routing

Open Shortest Path First (OSPF) is the next generation Internet routing protocol. You can configure the Pipeline 220 to communicate with other OSPF routers within an autonomous system (AS). To enable this routing function, you must configure the OSPF options on the Ethernet interface and for each WAN connection that supports remote OSPF routers.

OSPF can import routes from RIP as well. You can control the way these imported external routes are handled by adjusting system-wide routing options such as route preferences and ASE-type metrics.

Virtual Private Networking (VPN)

Many sites use the Internet to connect corporate sites or to enable mobile nodes to log into a corporate backbone. Such virtual private networks use cross-Internet tunneling to maintain security or to enable the Internet to transport protocols that it would otherwise drop, such as IPX. To implement virtual private networks, the Pipeline 220, with the VPN option, supports the Ascend Tunneling Management Protocol (ATMP) and the Point-to-Point Tunneling Protocol (PPTP).

ATMP enables the Pipeline 220 to create and tear down a tunnel to another Ascend unit. In effect, the tunnel collapses the Internet cloud and provides what looks like direct access to a home network. Packets received through the tunnel must be routed, so ATMP applies only to IP or IPX networks at this time.

Note: The Pipeline 220 supports PPTP by routing or forwarding PPTP traffic as appropriate. The Pipeline 220 does not act as either a PPTP Access Concentrator (PAC) or a PPTP Network Server (PNS).

Overview of management features

This section describes the following management functions, which use features built into the Pipeline 220:

- Using the Ascend Configurator
- Using the terminal server command line
- Using status windows to track WAN or Ethernet activity
- Managing the Pipeline 220 by means of SNMP
- Using remote management to configure far-end Ascend units
- Updating software in the Pipeline 220 unit's flash RAM

The Pipeline 220 provides up to nine security levels to control which management and configuration functions users can access.

Using the Ascend Configurator

To configure the Pipeline 220, you use the Ascend Configurator. This application is easily installed on your Windows NT workstation or Windows 95 workstation. The configurator enables you to:

- Configure the Pipeline 220 for the first time.
- Modify a pre-configured Pipeline 220.
- Save any Pipeline 220 configuration to a text file.

For more details, see Chapter 3, "Using the graphical interfaces." for more details.

Using the terminal server command line

To invoke the terminal-server command-line interface, you must use the VT100 interface and must have administrative privileges. Once you have activated a Security profile that enables the necessary privileges, you can invoke the command line by selecting Term Serv in the Sys

Diag menu. To close the command line, enter the Quit command at the command-line prompt. The cursor then returns to the VT100 menu interface.

Using status windows to track WAN or Ethernet activity

In the Pipeline 220 configuration menus, the right side of the screen displays eight status windows. The windows provide a great deal of read-only information about what is currently happening in the Pipeline 220. If you want to focus on the activity of a particular slot card, you can change the default contents of the windows to show what is currently going on in that slot.

Managing the Pipeline 220 by means of SNMP

Many sites use Simple Network Management Protocol (SNMP) applications to obtain information about the Pipeline 220. They use the information to enhance security, set alarms for certain conditions, and perform simple configuration tasks.

The Pipeline 220 supports the Ascend Enterprise MIB, MIB II, and some ancillary SNMP features. The Pipeline 220 can send management information to an SNMP manager without being polled. SNMP security uses a community name sent with each request. The Pipeline 220 supports two community names, one with read-only access, and the other with read/write access to the MIB.

Flash RAM and software updates

Flash RAM technology enables you to perform software upgrades in the field without opening the unit or changing memory chips. You can upgrade the Pipeline 220 through its serial control port, or you can use Trivial File Transfer Protocol (TFTP) to upgrade the unit through its Ethernet interface.

Where to go next

When you have planned your network, you are ready to configure the Pipeline 220. The flexibility of the Pipeline 220 and its ever-increasing number of configuration options means there is no set order for configuration. You can perform configuration tasks in any order you want. Figure 1-5 shows you where to look for the information you need. You can skip the chapters that do not apply to your installation.

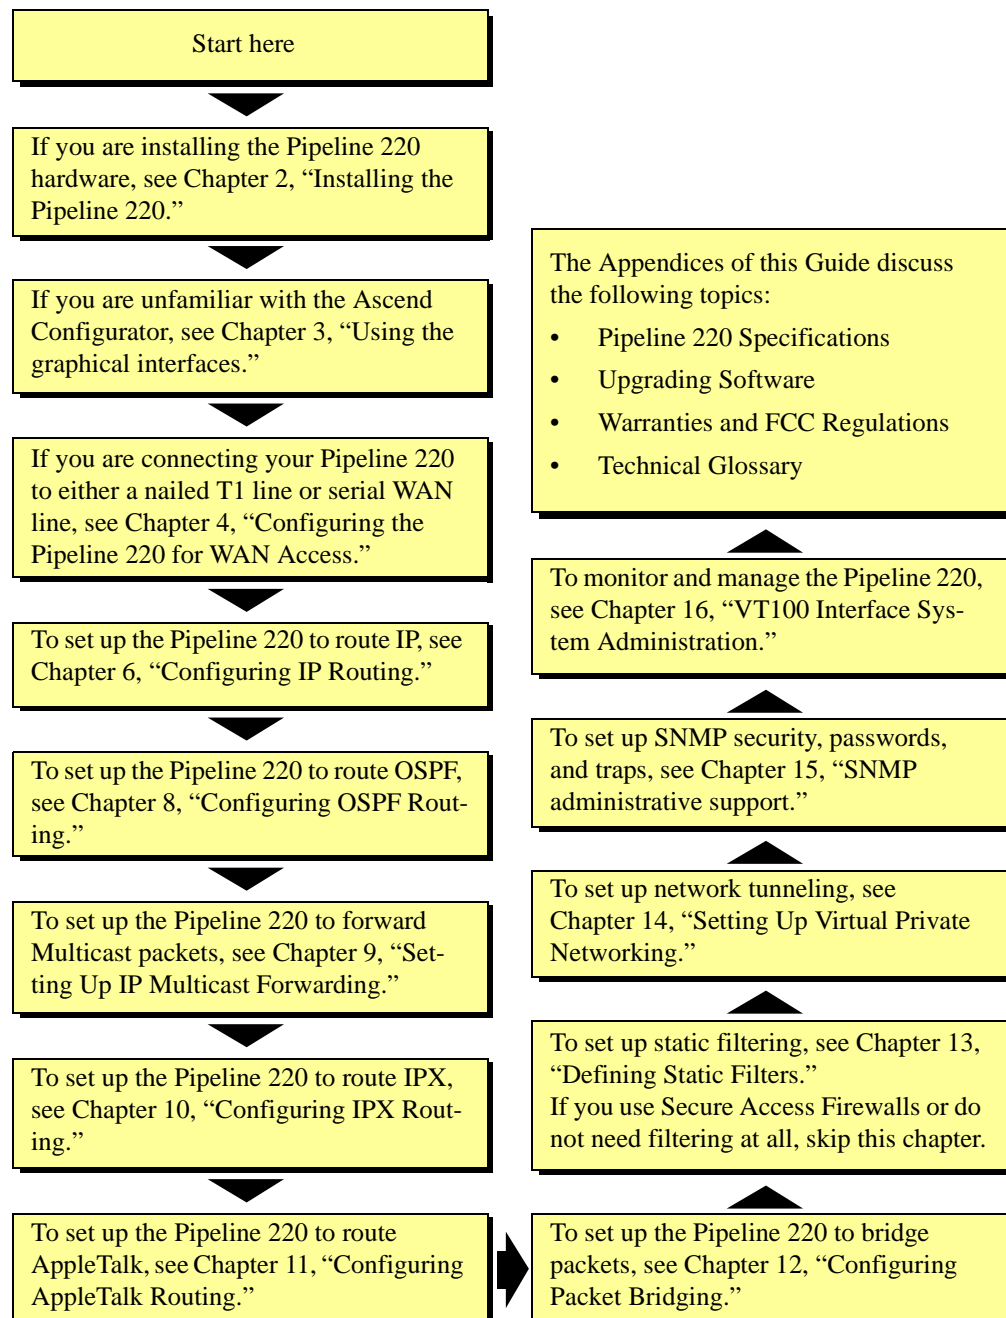


Figure 1-5. Pipeline 220 User's Guide Roadmap

Installing the Pipeline 220

This chapter contains:

What is included with the Pipeline 220?	2-1
Additional required hardware	2-2
Required software	2-3
Installation overview	2-4
Choosing a location for the Pipeline 220	2-5
Connecting the Pipeline 220 to the computer's Ethernet interface	2-5
Connecting to an Ethernet network	2-6
Connecting a computer to the Pipeline 220 Terminal port	2-6
Connecting the Pipeline 220 to your leased line	2-8
Starting up the Pipeline 220	2-8
Interpreting the Pipeline 220 LEDs	2-9

What is included with the Pipeline 220?

The Pipeline 220 package contains the following components:

- The Pipeline 220 unit.
- An RJ-48C T1/PRI crossover cable (part number 2510-0059-001).



Caution: This cable is for devices that transmit on pins 4 and 5 and receive on pins 1 and 2. Ask your carrier whether you should use a straight-through or crossover cable.

- A 10Base-T Ethernet crossover cable (part number 2510-0084-001).



Caution: If you are using the Pipeline 220 with only one computer and the computer has a 10Base-T Ethernet interface, you can use this cable to connect the computer directly to either of the 10Base-T Ethernet jacks on the Pipeline 220, as shown later in this chapter. *Do not use this cable for any other purpose.*

- A DB-9 to DB-25 serial-cable adapter (part number 2510-0052-002). You might need this adapter to connect the Terminal port on the Pipeline 220 to a serial port on your computer, as described later in this chapter.

- A power cable.

In addition to the items that came with your Pipeline 220, you must have additional hardware and software, which are explained in the next two sections.

Additional required hardware

In addition to the items supplied with the Pipeline 220, you need a WAN interface, a computer with a serial port, a modem cable, and an Ethernet interface.

WAN interface

Depending on the model of Pipeline 220 you have, you must have either a Serial WAN line or leased T1 line.

Computer with a serial port

To configure and monitor the Pipeline 220, you need a computer with a serial communication port capable of transmitting data at 9600 bits per second. The serial communication port is normally one you could use to connect an external modem. If you are not already familiar with your computer's serial ports, refer to your computer's user guide for more information.

If possible, you should set aside a serial port for a permanent connection to the Pipeline 220. Although a permanent connection with the serial port is not necessary for communication to the remote network, it allows you to monitor the Pipeline 220 at any time, manually connect to and disconnect from remote networks, and make configuration changes whenever necessary. Alternatively, once you have completed the initial configuration, you can communicate with the Pipeline 220 via Telnet.

Modem cable

To connect the Pipeline 220 to your computer's serial port, you need a modem cable (a serial communication cable designed for connecting an external modem). The cable must be a *high-speed* modem cable, that is, one that supports the *hardware handshaking* technique used by almost all recently manufactured modems. The cable, used in combination with the DB-9 to DB-25 adapter supplied with your Pipeline 220, must provide a 9-pin male D connector at one end and a plug that fits your available serial port at the other end.

Ethernet interface

For the Pipeline 220 to transmit data to and receive data from your computer, you need a properly configured 10Base-T (twisted-pair) or Thinnet Ethernet interface for your computer. The interface can be built into the computer, or it can be an adapter card in your PC or a PCMCIA card in your laptop. To install and/or configure the interface, follow the instructions included with the interface or with your computer.

Required software

To use your Pipeline 220, you need networking software and communications software.

Networking software

Depending on the type of network to which you will be connecting, you must have the appropriate networking software installed on your computer, as follows:

- If you are connecting to a Novell IPX network, you must have IPX client software.
- If you are connecting to the Internet or to a TCP/IP network, you must have software that supports TCP/IP networking. Many operating systems, such as Windows 95 and newer Macintosh OSes, include this networking software. If TCP/IP software is not included in your operating system, you need to obtain a separate software package.
- If you are connecting to an AppleTalk network, both the AppleTalk and TCP/IP software you need are included with newer Macintosh OSes.

Once you have networking software installed, you must configure it so that it can communicate with the Pipeline 220 and the remote network.

If you are unsure of the kind of software you must have installed on your computer, ask the network administrator or your Internet Service Provider (ISP).

Communications software

To configure and monitor the Pipeline 220, you need communications software that enables your computer to access the Pipeline 220 configuration interface. The software must be able to:

- Emulate a VT100 terminal.
- Communicate directly with the Pipeline 220, through the serial port to which the Pipeline 220 is connected rather than through a modem.

Most communications software that you purchase separately (such as the PROCOMM PLUS program for Microsoft Windows) works reliably. For Macintosh computers, a shareware communications program, ZTerm, works well.



Caution: The Terminal program included with Microsoft Windows 3.1 and the HyperTerm program included with Microsoft Windows 95 are not reliable enough for configuring the Pipeline 220.

Installation overview

When you install the Pipeline 220, you must perform the installation tasks in the order in which they are presented. The tasks are described in the following sections:

- 1 “Choosing a location for the Pipeline 220” on page 2-5.
- 2 “Connecting the Pipeline 220 to the computer’s Ethernet interface” on page 2-5
or
“Connecting to an Ethernet network” on page 2-6.
- 3 “Connecting a computer to the Pipeline 220 Terminal port” on page 2-6.
- 4 “Connecting the Pipeline 220 to your leased line” on page 2-8.
- 5 “Starting up the Pipeline 220” on page 2-8.

Once you have successfully installed the Pipeline 220 you can proceed to configure it as explained in Chapter 3, “Using the graphical interfaces.”

Figure 2-1 illustrates the Pipeline 220 backpanel on which the unit’s connectors are located.




Figure 2-1. Back panel of the Pipeline 220

Table 2-1 describes the elements of the Pipeline 220 backpanel.

Table 2-1. Pipeline 220 ports

Element	Description
Power supply	Connects to the power supply. For power requirements, refer to Appendix A, “Pipeline 220 Specifications. ”
Alarm	Alarm relay. For Alarm relay specifications, refer to Appendix A, “Pipeline 220 Specifications. ”
Control	Connect a serial cable here to access the Pipeline 220 VT-100 configuration interface.
WAN 2	Connects via a V.35 cable to a DCE device. Note that if your Pipeline 220 has a nailed T1 line, this port is not supported.
WAN 1	Connects via an RJ-45 cable to a nailed T1 line. Note that if your Pipeline 220 has a serial WAN port, this port is not supported.
PCMCIA	To be used in future releases.
DRAM	To add additional DRAM memory.

Table 2-1. Pipeline 220 ports (continued)

Element	Description
ENET 1	Connects via 10Base-T cable to an Ethernet segment.
ENET 2	Connects via 10Base-T cable to an Ethernet segment.

Choosing a location for the Pipeline 220

If possible, place the Pipeline 220 in a location that lets you view the lights on the front. These lights show the current status of the Pipeline 220, such as whether the WAN line or Ethernet lines are in use, and can help you diagnose problems.

Connecting the Pipeline 220 to the computer's Ethernet interface

If you connect only one computer to a particular Ethernet port on the Pipeline 220, and the computer has a 10Base-T (twisted-pair) Ethernet interface, you can connect the Pipeline 220 and computer with the special 10Base-T cable, known as a *crossover cable*, which is included with the Pipeline 220. Proceed as follows:

- 1 Insert one end of the 10Base-T crossover cable (part number 2510-0084-001) into one of the 10BT jacks on the back of the Pipeline 220.
- 2 Insert the other end of the cable into the 10Base-T Ethernet jack on the computer.

Proceed to “Connecting a computer to the Pipeline 220 Terminal port” on page 2-6.

Connecting to an Ethernet network

The Pipeline 220 has two separate Ethernet interfaces, and you must configure each interface separately. If you configure both interfaces, the Pipeline 220 routes, bridges, or routes and bridges between the two interfaces.

To connect the Pipeline 220 to a 10Base-T hub:

- 1 Insert one end of a 10Base-T cable into one of the 10BT jacks on the back of the Pipeline 220.



Caution: Do not use the 10Base-T crossover cable included with the Pipeline 220 (part number 2510-0084-001) to connect the Pipeline 220 to a 10Base-T hub. This cable is only for connecting the Pipeline 220 directly to a computer, as described in “Connecting the Pipeline 220 to the computer’s Ethernet interface” on page 2-5.

- 2 Insert the other end of the cable into an unused port on the 10Base-T hub.

Connecting a computer to the Pipeline 220 Terminal port

To access the Pipeline 220’s VT100 interface, you use a computer and a serial communications connection as described in Chapter 16, “VT100 Interface System Administration.” Use the VT100 Interface to access the Pipeline 220 Terminal Server interface and status windows. The following sections explain how to connect different types of computers to the Pipeline 220:

- If you are using an IBM-compatible personal computer to configure the Pipeline 220, see “Connecting an IBM-compatible computer” on page 2-6.
- If you are using a Macintosh computer to configure the Pipeline 220, see “Connecting a Macintosh” on page 2-7.
- If you are using a Unix workstation to configure the Pipeline 220, see “Connecting a Unix workstation” on page 2-7.

Connecting an IBM-compatible computer

To connect an IBM-compatible personal computer to the Pipeline 220:

- 1 Find an unused serial connector on your computer.
Make a note of which serial port you are connecting to (most often COM1 or COM2). You will need this information later when setting up the configuration software. If no serial port is currently free, disconnect from one of the ports a device that you can temporarily do without, such as an external modem.
- 2 Connect a modem cable to the serial connector.
- 3 If the plug at the other end of the modem cable has 25 pins, connect it to the 25-to-9 pin adapter included with the Pipeline 220 (part number 2510-0052-002).
- 4 Connect the cable to the Terminal port on the back of the Pipeline 220.

Proceed to “Connecting the Pipeline 220 to your leased line” on page 2-8.

Connecting a Macintosh

To connect a Macintosh or compatible personal computer to the Pipeline 220:

- 1 Connect the 25-to-9 pin adapter included with the Pipeline 220 (part number 2510-0052-002) to the DB-25 end of a Macintosh modem cable.




Figure 2-2. Connecting an adapter to a Macintosh modem cable

- 2 Connect the cable to the Terminal port on the back of the Pipeline 220.
- 3 Connect the other end of the cable to a serial port (either the Modem or Printer port) on the computer.

Proceed to “Connecting the Pipeline 220 to your leased line” on page 2-8.

Connecting a Unix workstation

To connect a workstation or other computer running Unix to the Pipeline 220:

- 1 Connect a modem cable for the computer to the Terminal port on the back of the Pipeline 220.
- 2 Connect the other end of the cable to the serial port on the computer.

Proceed to “Connecting the Pipeline 220 to your leased line” on page 2-8.

Connecting the Pipeline 220 to your leased line

Before you can configure the Pipeline 220, you have to connect it to your leased line. Before connecting the unit to your leased line, you must obtain the carrier's approval.



Caution: To avoid harming the WAN, you must contact your carrier for approval before installation. Once you have installed the Pipeline 220, you must notify the carrier before disconnecting the Pipeline 220 from the WAN. If you disconnect or turn off the Pipeline 220 without prior notification, the carrier might temporarily discontinue your leased service.

To connect the Pipeline 220 to the leased line that provides access to your WAN, proceed as follows:

- 1 Connect your leased line to the associated WAN port on the Pipeline 220. Connect the other end of the leased line either directly or through other network interface equipment.

Note: For connection to the demarcation point, where the T1 line's metallic interface connects to other equipment, the Pipeline 220 T1 ports are equipped with internal CSUs.

- 2 Inform your service provider that your equipment is connected, so the provider can bring up the line.

Starting up the Pipeline 220

When you boot the Pipeline 220, position yourself so you can watch the LEDs on the front panel.

After you power up the Pipeline 220, it takes less than a minute for it to become ready to use. The status light labeled PWR on the front of the Pipeline 220 comes on immediately to indicate that the power is on.

The LEDs labeled DRAM card, PCMCIA card, WAN Act, and Alarm illuminate when you connect the Pipeline 220 to its power source, and go off when the unit passes the internal Power-On Self Test (POST). After successful POST, the Pipeline 220 is ready for you to use. Chapter 3, "Using the graphical interfaces," describes how to perform basic configuration of your Pipeline 220.

If there is a problem with the card slots or WAN interface, their associated LEDs flash. Refer to Figure 2-2 for descriptions of all the Pipeline 220 LEDs.

Interpreting the Pipeline 220 LEDs

Figure 2-3 illustrates the Pipeline 220 front panel and location of LEDs.




Figure 2-3. Front panel of the Pipeline 220

Table 2-2 explains the Pipeline 220 LEDs.

Table 2-2. Pipeline 220 LEDs

LED	Indicates (when lit)
Ethernet ACT 2	Activity on the Ethernet interface.
Ethernet ACT 1	Activity on the Ethernet interface.
DRAM Card	Activity on the DRAM card.
PCMCIA Card	To be used in future releases.
WAN Act	Activity on the WAN interface.
Power	Unit is powered on.
Alarm	WAN alarm or a trunk is out of service, such as during line loopback diagnostics. WAN alarms include Loss of Sync, Red Alarm, Yellow Alarm, and All Ones (or AIS).

The next chapter explains how to turn on the Pipeline 220 and use the graphical interfaces to configure it.

Using the graphical interfaces

This chapter explains how to use the Ascend Configurator interface and the Quickstart interface. It includes the following sections:

Before you begin	3-1
Installing the Ascend Configurator	3-2
Using the Configurator Startup screen.....	3-3
Using the Ascend Configurator interface.....	3-7
Using the Pipeline 220 Quickstart	3-14

The Configurator provides a graphical interface for configuring your Pipeline 220. Quickstart prompts you for the information needed for a minimal configuration. To get your unconfigured Pipeline 220 up and running as quickly as possible, first assign it an IP address, then access Quickstart and supply settings through the supplied sequence of screens.

If you take this approach, you should use the Configurator again later to optimize your configuration. For example, you might use the information in this chapter to assign an IP address to the Pipeline 220, but you should also read Chapter 6, “Configuring IP Routing,” for a complete discussion of IP routing and the steps you follow to configure all IP-routing features.

This chapter does not explain which values to choose when entering the configuration settings. But the Quickstart section includes cross references to the chapters that explain the settings you are prompted for, and the Configurator Online Help provides detailed information about all the settings.

Before you begin

Before installing and using the Ascend Configurator, make sure you have installed the Pipeline 220 as explained in the previous chapter. If you have not already assigned the Pipeline 220 unit an IP address, the unit must be on the same physical Ethernet as the workstation running the Ascend Configurator. The Ascend Configurator will not find a completely unconfigured Pipeline 220 units across a router. (The unit must first have a valid IP address.)

Also make sure that you have:

- Access to a workstation running Windows NT or Windows 95.
- The Telnet password for the Pipeline 220 you want to configure. The default is a null.
- The SNMP read/write community string. The default is *write*.

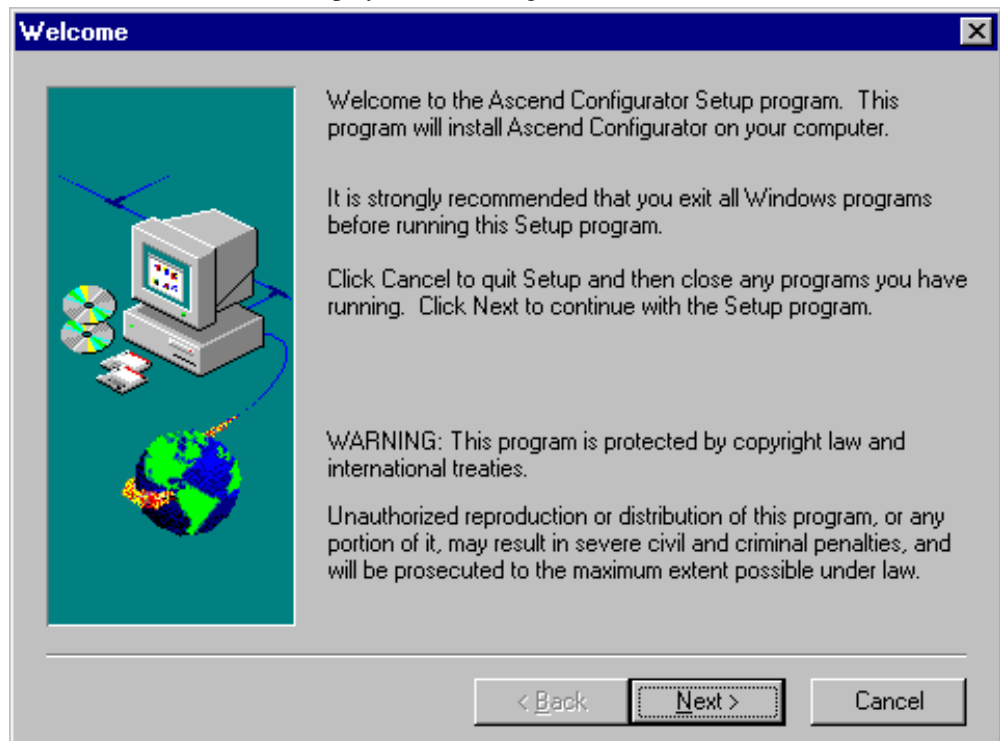
- An HTML-capable browser installed on the workstation to access the Configurator Online help.

Installing the Ascend Configurator

To install the Ascend Configurator, copy the Configurator executable file to your Windows NT or Windows 95 workstation.

Click the Windows Start button, then select Run. Use the Browse button to locate the executable file, and click OK to run it.

A Wizard loads the file and displays the following screen:



This screen is the first of several that explain the installation process and prompt you for additional information. In most circumstances, you should accept the default values.

Refer to the README file for additional information about installing the Ascend Configurator.

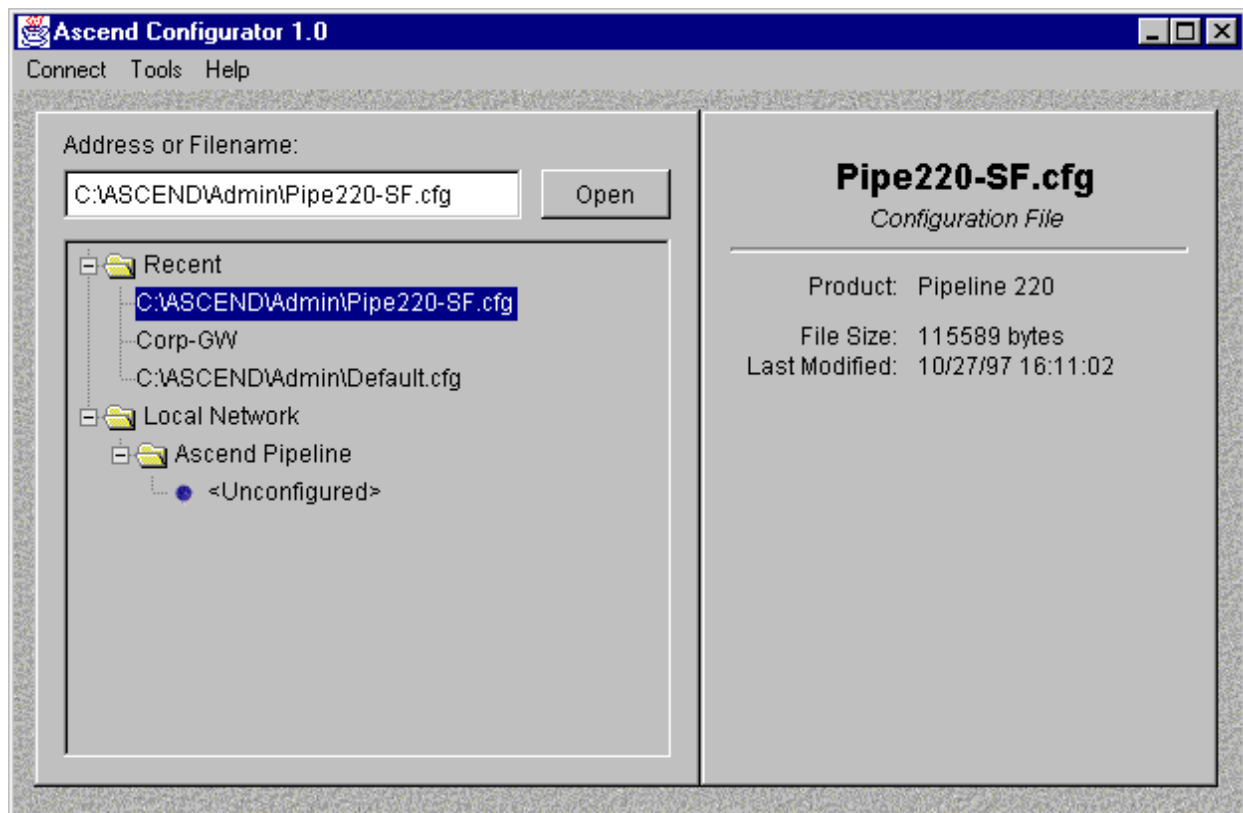
Using the Configurator Startup screen

The Configurator's Startup screen enables you to:

- Connect to any unconfigured Pipeline 220 on the same IP subnet as the PC on which the Configurator is running.
- Connect to any configured Pipeline 220 accessible on your IP network.
- Set optional features of the Configurator.

Settings on the Startup screen

After you launch the Ascend Configurator, the Startup screen appears:



The Startup screen contains the following elements:

Element	How to use it
Address or Filename Entry Field:	Specify the address or filename of the Ascend product you want to configure.
Open button:	Click Open to connect to the specified address or configuration filename.
Recent folder:	Highlight an item listed in the Recent folder to open a previously accessed address or configuration file.

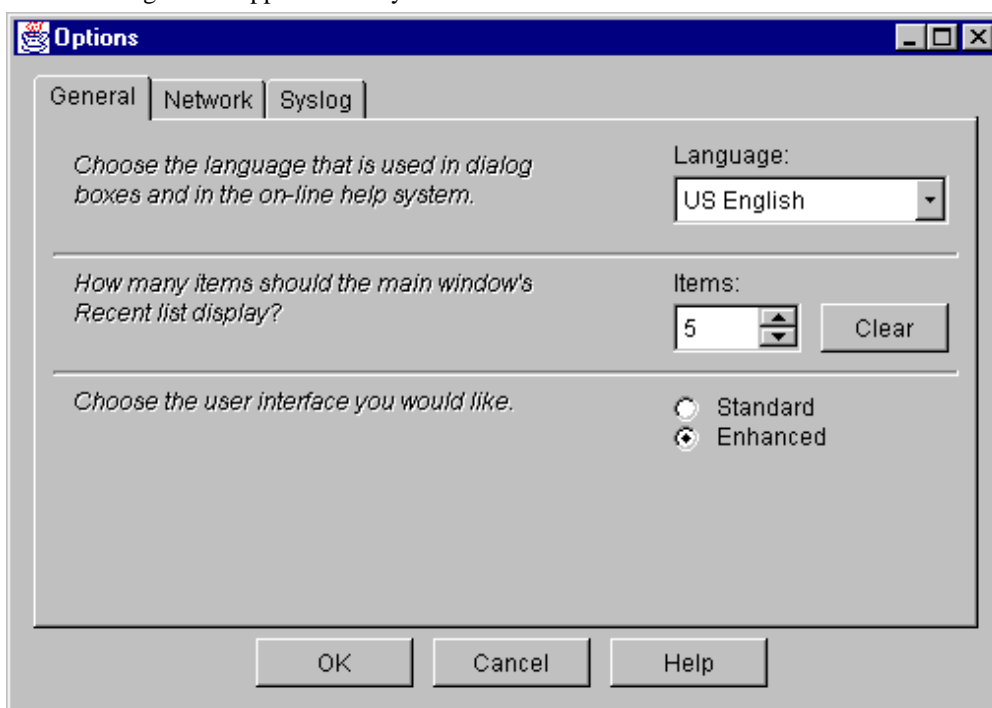
Element	How to use it
Tools menu	Open the Tools menu, then select Options to access the Options screen.

Using the Options screen

You can access the Options screens via the Tools drop-down menu. The Options screen displays three tabs: General, Network, and Syslog.

Settings on the General tab screen

The following screen appears when you click the General tab:

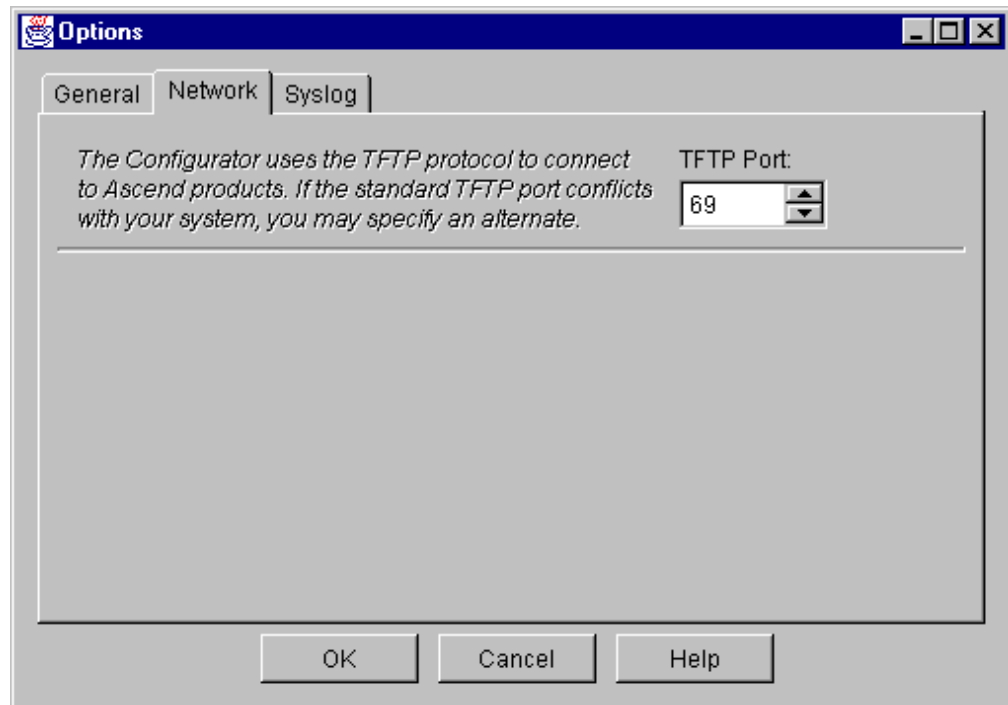


Configure the following General settings:

Setting	Description
Language:	Enables you to choose the language used in the dialog boxes and help screens.
Recent Display:	Enables you to choose the number of recent items displayed in the main screen.
User Interface:	Allows you to choose between a standard or an enhanced display.

Settings on the Network tab screen

The following screen appears when you click the Network tab:



Configure the following Network setting:

Setting	Description
TFTP Port:	Specifies the port used for TFTP access to your Ascend product.

Settings on the Syslog tab screen

The following screen appears when you click the Syslog tab:

The screenshot shows a Windows-style dialog box titled "Options". It has three tabs: "General", "Network", and "Syslog", with "Syslog" being the active tab. The dialog contains the following elements:

- A text area with the message: *The Configurator listens for Syslog messages on the standard Syslog port. If this port conflicts with your system, you may specify an alternate.*
- A "Syslog Port:" label next to a spin box containing the value "514".
- A text area with the message: *The Configurator can keep a history file of all Syslog activity. How would you the Configurator to manage the history file?*
- Three radio buttons for managing the history file:
 - ☒ Disabled
 - ☐ Append messages
 - ☐ Overwrite messages
- A text area with the message: *Specify the name of the history file:*
- A text box containing the value "Syslog".
- A checkbox labeled "Disable Syslog" which is currently unchecked.
- At the bottom, there are three buttons: "OK", "Cancel", and "Help".

Configure the following Syslog settings:

Setting	Description
Syslog Port:	Specifies the UDP port on which the Configurator should listen for Syslog messages.
History file	Specifies how the Configurator's Syslog history file is maintained.
History file name	Specifies the name of the Syslog history file
Disable Syslog	Disables Syslog on the Configurator.

Using the Ascend Configurator interface

The Ascend Configurator interface (Figure 3-1) consists of a series of tabs that you use to configure your Pipeline 220. The interface is dynamically created, depending on the hardware configuration of your Pipeline 220.

Note: Screen size limits the number of tabs that can be displayed at one time. Use the scroll buttons, to the right of the tabs, to access additional tab categories.



Figure 3-1. Ascend Configurator interface

It is important to understand that while you are using the Ascend Configurator you are not connected to the Pipeline 220. That is, none of the changes you make to the Pipeline 220 configuration take effect until you re-establish a connection to the Pipeline 220 and upload the changed configuration. You do this with the Upload command, as explained in "Uploading a configuration" on page 3-13.

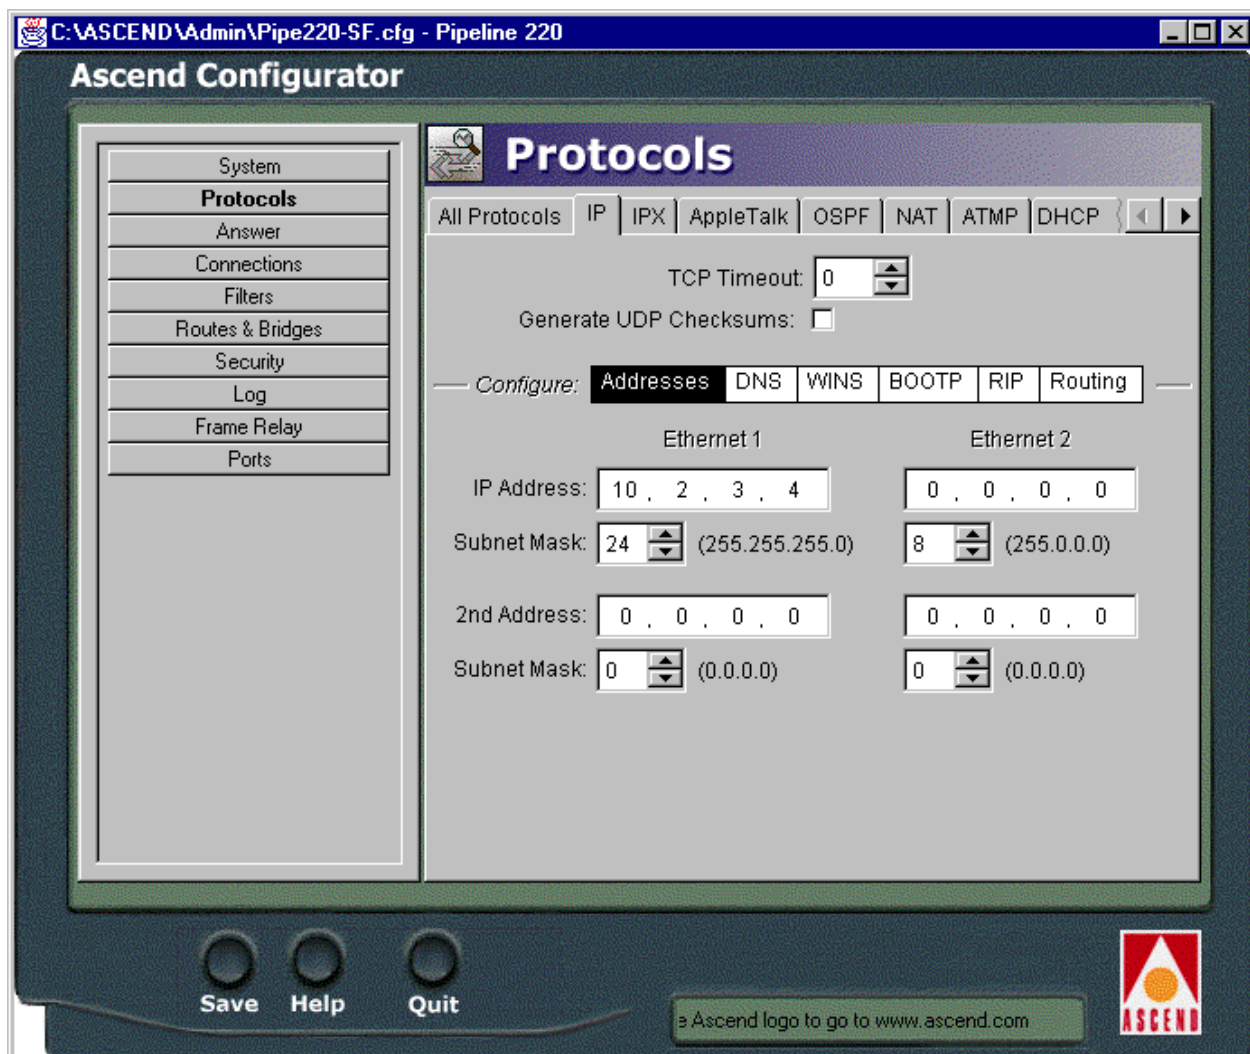
Using the graphical interfaces

Using the Ascend Configurator interface

In some cases, this manual refers to Configurator parameters by indicating the sequence of button and tab clicks required to access the parameter. For example, the steps you must perform to access the IP Address parameter are as follows:

- 1 On the left side of the Configurator, click the Protocols button.
- 2 On the right side of the Configurator, click the IP tab.
- 3 On the lower-right section of the Configurator, click the Addresses button.

The IP Address appears in the lower-right section of the Configurator:



This manual might refer to the parameter (also indicating the sequence of steps required to access the parameter) as the Protocol button > IP tab > Addresses button > IP Address parameter.

Ascend Configurator tabs

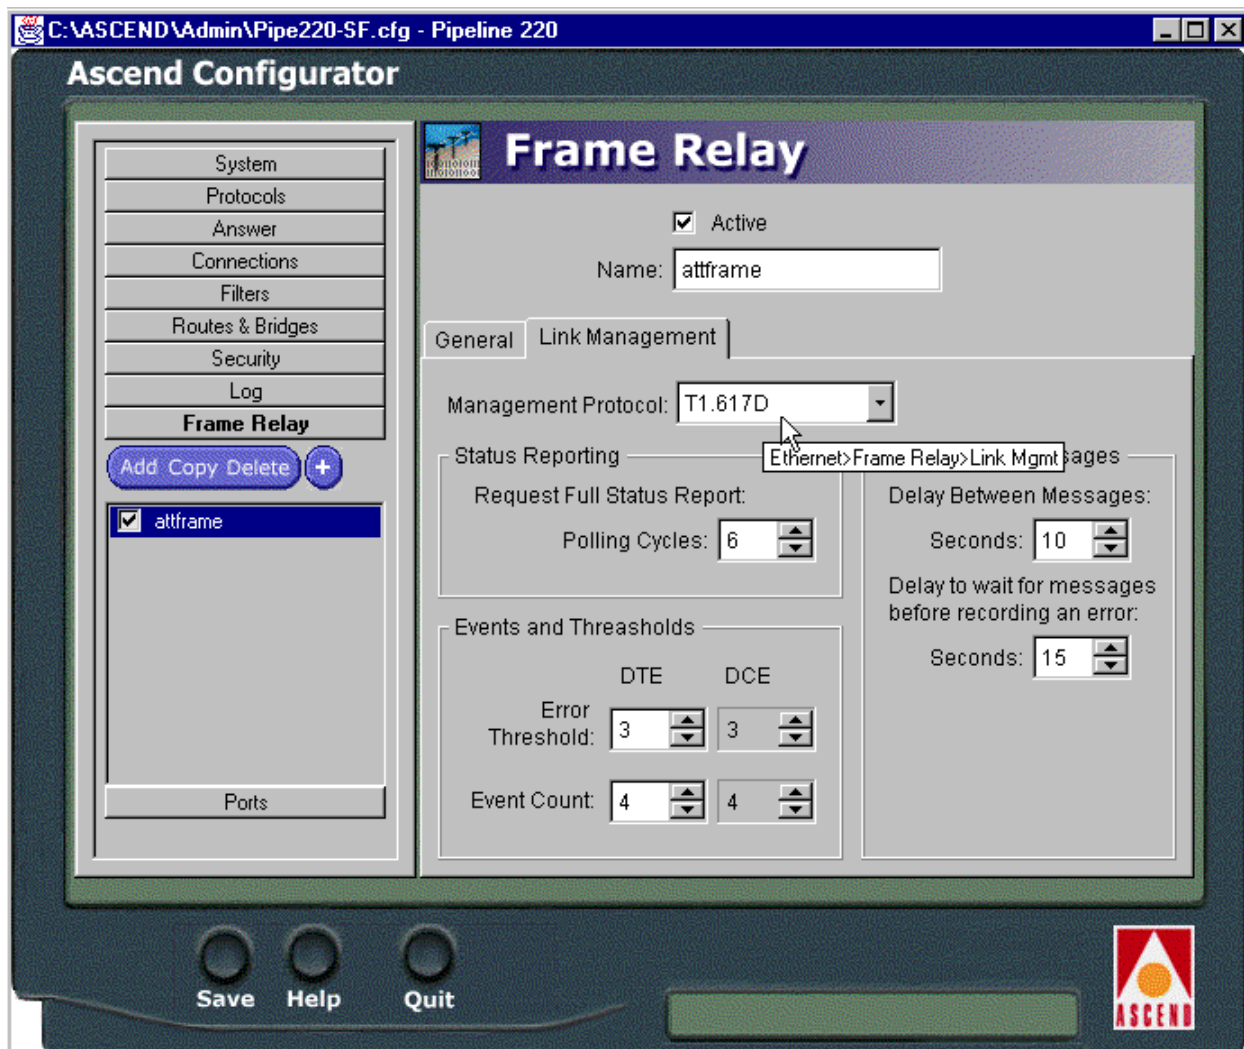
The Ascend Configurator interface consists of the following tabs:

Tab name	Description
System	Defines system level settings
Protocols	Displays settings that apply to all protocols on the Pipeline.
Answer Options	Defines the Pipeline 220's default settings when answering calls, including terminal server calls. More complicated connections require that you define a Connection profile.
Connections	Defines profiles, which are used by remote devices to establish a connection to the Pipeline 220, or by the Pipeline 220 to establish a connection with a remote device.
Filters	Defines static profiles for Generic, IP, or IPX input and output filters. Apply them to specific connections, to the default Answer profile, or to the Ethernet interface.
Routes and Bridges	Defines static configuration profiles for the Pipeline 220's routing or bridging tables.
Security	Defines security profiles, thereby allowing various levels of access to the Pipeline 220.
Log	Enables the Pipeline 220 to send Syslog messages to a Syslog host.
Frame Relay	Defines parameters used for Frame Relay connections.
Ports	Enables either the T1-CSU port or the Serial WAN port, and defines parameters used for the connection.

If you are familiar with Ascend's VT100 user interface

For every Configurator parameter, there is an equivalent VT100 user interface parameter. The Ascend Configurator displays the equivalent VT100 parameter if you position the mouse pointer on the field of any Configurator parameter. You must leave the pointer stationary over the field for two seconds.

Positioning the pointer over the Link Management protocol field shows its VT100 equivalent, Link Mgmt:



Assigning the Pipeline 220 an IP address

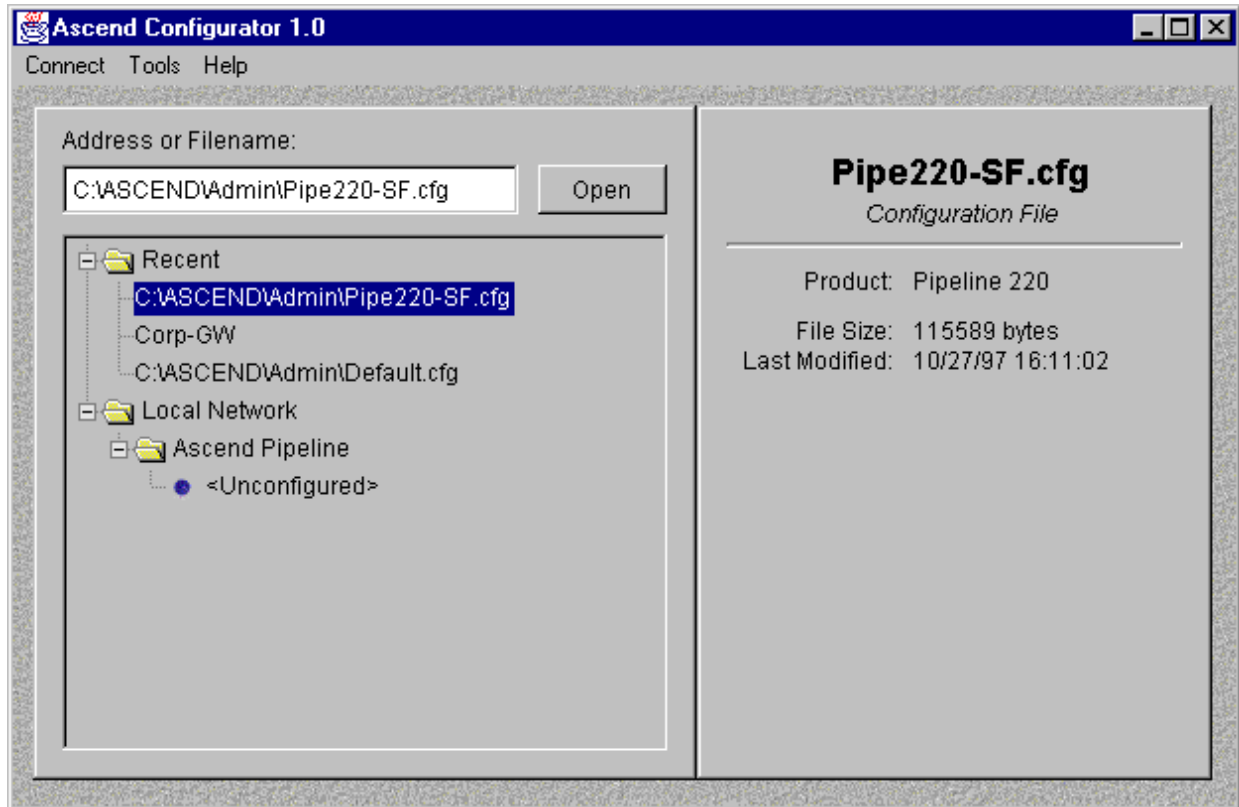
Before you assign the Pipeline 220 an IP address, make sure that:

- The Ascend Configurator is running on a workstation that resides on the same physical subnet as the unconfigured Pipeline 220.

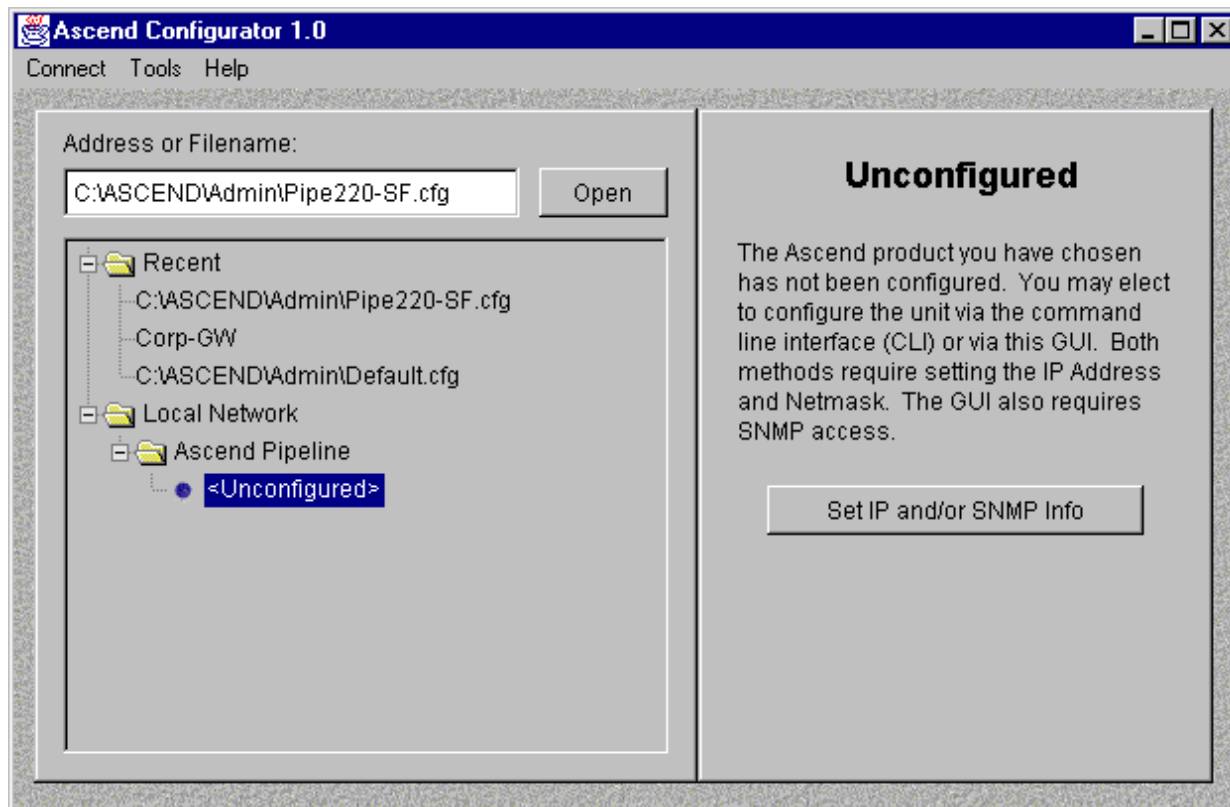
Note: To connect to a Pipeline 220 with an assigned IP address, the workstation does *not* have to reside on the same physical segment as the Pipeline 220.

- The Pipeline 220 is powered up, has passed its Power-On Self Test, and is connected to the Ethernet network.

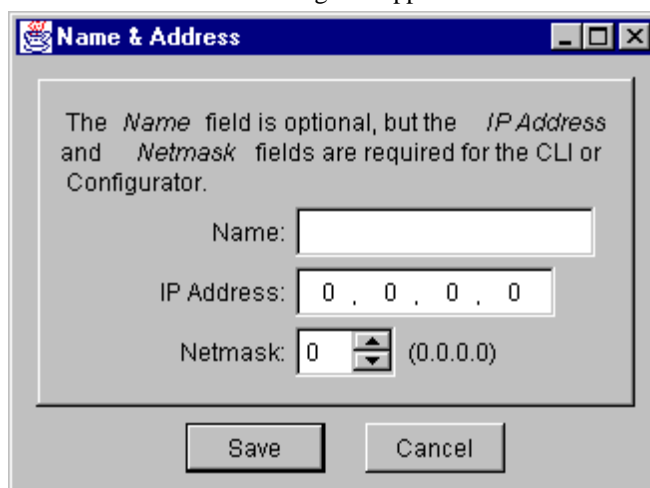
To assign the Pipeline 220 an IP address, first launch the Configurator. The Connect screen appears:



- 1 Double-click the label <Unconfigured>. The Set IP button appears on the right side of the Configurator:



- 2 Click the “Set IP and/or SNMP Info” button.
The Name and Address dialog box appears:

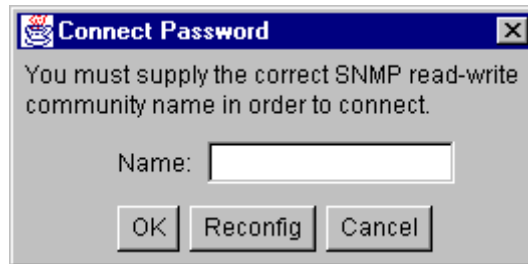


- 3 Assign an IP address, subnet mask, and, if necessary, a name.
- 4 Click the Save button.
The Ascend Configurator connects to the Pipeline, assigns it the address you have specified, and downloads its configuration via TFTP. After a few moments, the configuration interface appears.

Connecting to a Pipeline 220

You can connect to a Pipeline 220 once it has an IP address. To connect:

- 1 From the Connect screen, double-click the Name or IP address of the Pipeline 220 to which you want to connect. You can also enter the name or IP address of the Pipeline 220 in the “Address or Filename” box, and click the Open button.
- 2 If the following dialog box appears, enter the SNMP Read/Write community name.



The default value is write.

The Ascend Configurator connects to the Pipeline 220 and downloads its configuration via TFTP. After a few moments, the configuration interface appears.

Opening an existing configuration file

You can open an existing Pipeline 220 configuration file, that is stored in a local file system, and edit it using the Ascend Configurator. You can then upload the new file to any Pipeline 220, or save it as a configuration file in the local file system.

To open an existing Pipeline 220 configuration file:

- 1 From the Startup screen select Browse > Files.
The Open File dialog box appears.
- 2 Select the configuration file you want to edit.
The Ascend Configurator reads the configuration file. After a few moments, the configuration interface appears.

Uploading a configuration

The Pipeline 220 settings you configure with the Ascend Configurator do not take effect until you upload them to the Pipeline 220.

To upload the Pipeline 220 configuration file:

- 1 From the Ascend Configurator tab view, select Save.
- 2 If you are uploading the configuration to the same device, select Upload Changes to *IP Address*, where *IP address* is the IP Address of the Pipeline 220 to which you want to upload the configuration.
- 3 If you are uploading the configuration to a different device, select Upload to an Ascend Product and enter the IP address of the device.
- 4 Click Save.

The Ascend Configurator uploads the configuration file via TFTP.

Saving a configuration to a file

After you configure Pipeline 220 settings with the Ascend Configurator, you can save the settings to a file as a backup or for upload to Pipeline 220 units at a later time.

To save the Pipeline 220 configuration file:

- 1** From the Ascend Configurator tab view, select Save.
- 2** Select Save a Copy Under a New File Name.
- 3** Enter the name of the file, or click Browse if you want to overwrite an existing file.
- 4** Click Save.

Security issues with the Configurator

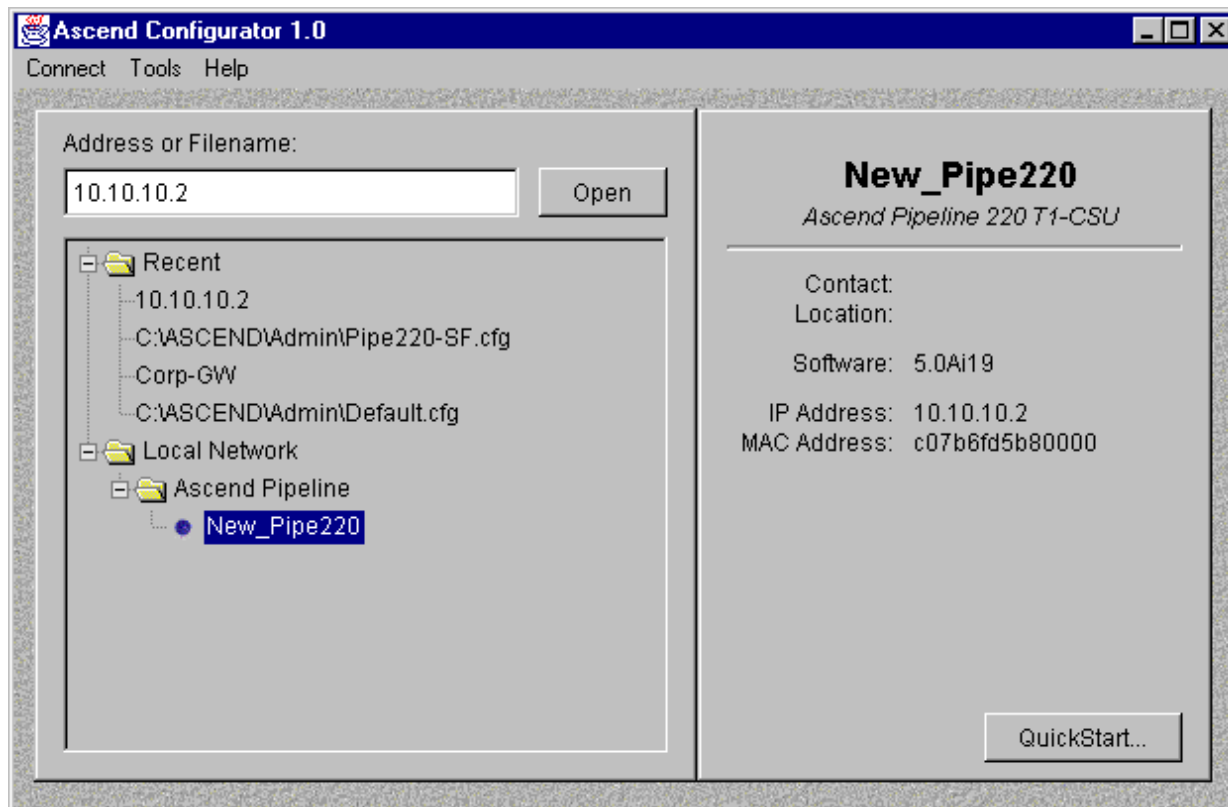
When using the Pipeline 220, take the following precautions to safeguard your Ascend unit from unauthorized access:

- Because the Ascend Configurator relies on SNMP to configure the Ascend unit, you should restrict the number of SNMP Managers allowed read/write access.
- Change the default read/write community string for the Ascend unit.
- Change the Telnet password.
- Communication with the Ascend unit is via TFTP. All configuration information (including passwords) are sent unencrypted. Therefore make sure that you run the Ascend Configurator only on a secure, trusted network

Using the Pipeline 220 Quickstart

Once you have supplied the Pipeline 220 with an IP address, the easiest way to configure your Pipeline 220 is by using the Quickstart screens. They enable you to perform basic Pipeline 220 configuration. Subsequently, you can use the Ascend Configurator to fine-tune your configuration.

When you connect to a Pipeline 220 with the Configurator, the lower-right section of the Connect screen displays the Quickstart button:




Click the Quickstart button to launch the Quickstart interface. The interface includes several configuration screens, which prompt you for the minimal configuration necessary to enable required protocols. Each screen includes a Previous and Next button for navigation between screens. The Quickstart screens are as follows:

Screen name	Description
System Information	Defines information identifying the Pipeline 220.
Internet Protocol (IP)	Specifies parameters necessary for IP communication. For complete information about IP, see Chapter 6, "Configuring IP Routing."
Encapsulation (Answer) Options	Specifies the WAN-encapsulation protocols the Pipeline 220 supports. For complete information about encapsulation parameters, see the Configurator Online Help.
IPX and AppleTalk Protocol	Allows you to enable either IPX routing or AppleTalk routing, or both. Use this screen to configure the necessary parameters. For complete information about IPX, see Chapter 10, "Configuring IPX Routing." For complete information about AppleTalk, see Chapter 11, "Configuring AppleTalk Routing."

Screen name	Description
Network Address Translation (NAT)	Allows you to enable NAT and configure the necessary parameters. For complete information about NAT, see Chapter 7, “IP Address Management.”
Remote Management	Defines a Telnet password needed for access to the Pipeline 220.
WAN Interface Configuration	Specifies the active Pipeline 220 WAN interface, and enables you to configure the necessary parameters. For complete information about WAN interfaces, see Chapter 4, “Configuring the Pipeline 220 for WAN Access.” For complete information about Frame Relay, see Chapter 5, “Configuring Frame Relay.”

Use the Previous and Next buttons on the Quickstart to navigate between screens.

When you click the Next button on the WAN Interface Configuration screen, the following dialog box appears:



Click Yes to upload the configuration to the Pipeline 220. Click No to return to the WAN Interface Configuration screen without uploading the configuration.

Click No to exit the Quickstart. If you have set parameters in the Quickstart, but have not uploaded them to the Pipeline 220, they are discarded.

Configuring the Pipeline 220 for WAN Access

This chapter covers the following topics:

Introduction to WAN configuration	4-1
Configuring the T1 line	4-1
Configuring the serial WAN port	4-4

Introduction to WAN configuration

The Pipeline 220 comes with a built-in T1-line connection and a V.35 serial port. You must configure one or the other for WAN access.

Figure 4-1 shows the Pipeline 220 back panel. The T1 connection is the RJ-45 port labeled WAN1. The V.35 serial port is labeled WAN2.




Figure 4-1. Back panel of the Pipeline 220

Configuring the T1 line

The T1 connection (WAN1) on the Pipeline 220 is not channelized, but you can configure it like a T1 with any number of DS0 channels (up to 24), as specified by your carrier.

With a nailed T1 line, you must manually configure some port information. For example, you must specify the signals that indicate that the Data Communications Equipment (DCE) is ready to connect. In addition, you might need to adjust the amount of attenuation that the Pipeline 220 should apply to the line's network interface in order to match the cable length from the Pipeline 220 to the next repeater.

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

- Supply information, such as encoding, framing, and buildout (attenuation) that you obtain from your carrier.
- Activate the port.

T1 parameters

This section provides background information about the T1 line-interface parameters.

For complete information about each parameter, see the *Configurator Online Help*.

T1 line framing and encoding

The framing used by the physical layer of the T1 line may be D4 or ESF. D4, also known as the superframe format, consists of 12 consecutive frames separated by framing bits. The line must not use ISDN signaling with D4 framing, because false framing and Yellow Alarm emulation can result. ESF specifies the extended superframe format. This format consists of 24 consecutive frames separated by framing bits. The ISDN specification advises that you use ESF with ISDN D-channel signaling.

The encoding value sets the layer-1 line encoding used for the physical links, which affects the way data is represented by the digital signals on the line. Your carrier can tell you which encoding to use. AMI (the default) specifies Alternate Mark Inversion encoding. B8ZS specifies that the encoding is Bipolar with 8-Zero Substitution. The other option, None, is identical to AMI but without density enforcement.

Amount of attenuation required

The Buildout parameter specifies the amount of attenuation to apply to the T1 transceiver's internal Channel Service Unit (CSU). The amount depends on the cable length from the Pipeline 220 to the next repeater. Valid values are 0 db (decibels) through 22.5 db.

Attenuation is a measure of the power lost on a transmission line or on a portion of that line. When you specify a value for Buildout, the Pipeline 220 applies attenuation to the T1 line, causing the line to lose power. Repeaters boost the signal on a T1 line. If the Pipeline 220 is too close to a repeater, you might need to add some attenuation. Check with your carrier to determine the correct value.

Clock source for synchronous transmission

The Clock Source parameter indicates whether the T1 line can be used as the master clock source for synchronous connections. In synchronous transmission, both the sending device and the receiving device must maintain synchronization in order to determine where one block of data ends and the next begins.

You might need to disable this parameter on one unit if two Ascend units connect to each other by a crossover cable (with optional T1 repeaters) between their network ports.

Configuring the nailed T1 line

To configure the nailed T1 line., you must enter the settings, then enable them.

Entering the settings

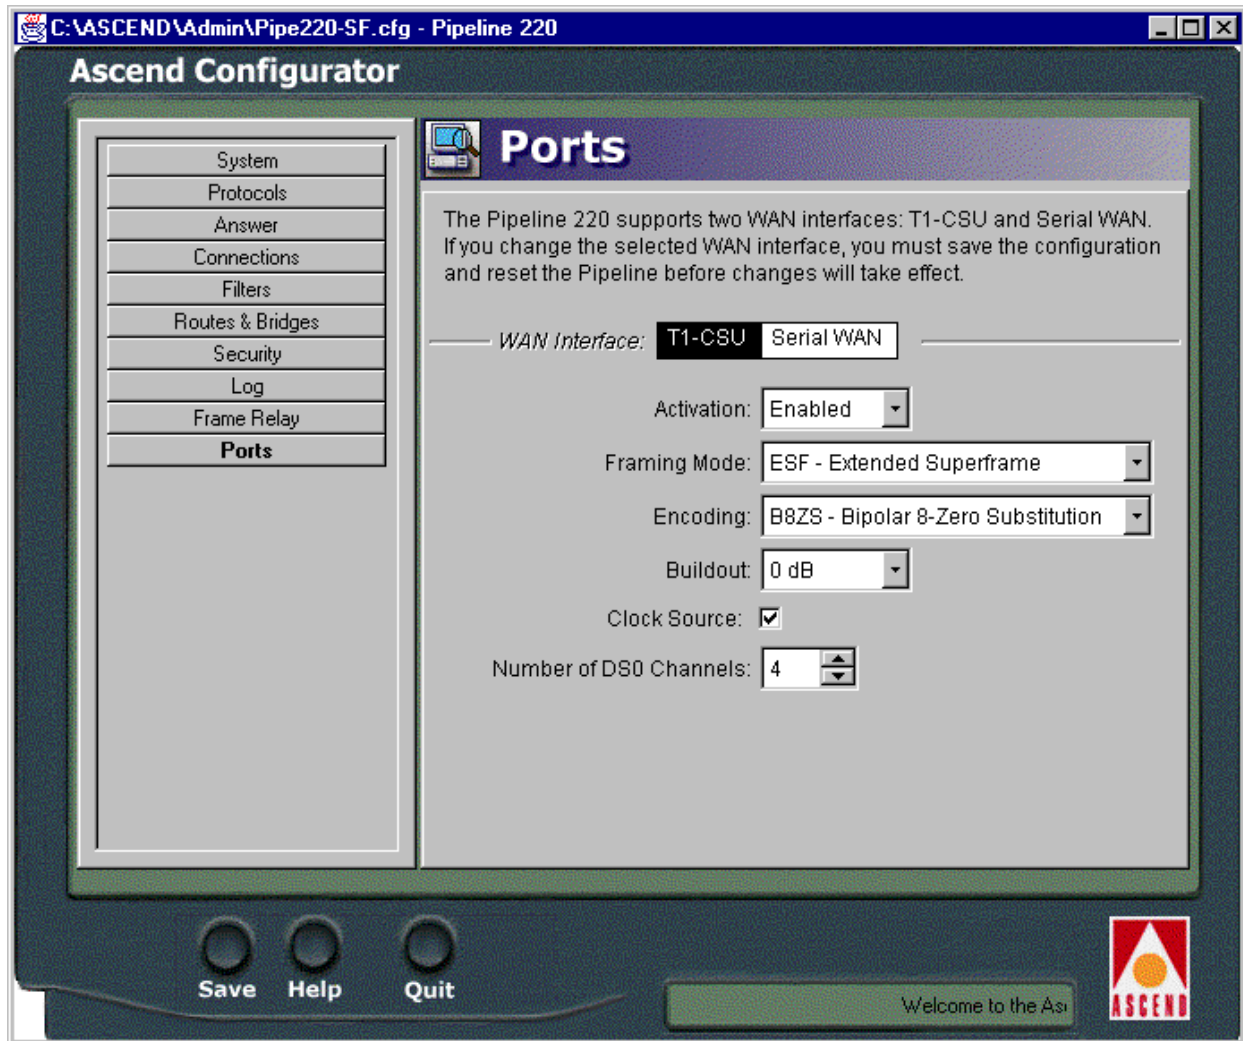
To enter the configuration of the nailed T1 line:

- 1 On the left side of the Configurator, click the Ports button.

The right side of the Configurator now displays WAN Interface buttons.

- 2 Click the T1-CSU button.

The Configurator now displays parameters specific to the nailed T1 interface:



- 3 Set Activation to Enabled.
- 4 Set the Framing Mode as specified by your carrier.
- 5 Set the Encoding parameter as specified by your carrier.
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.
- 6 Set the buildout if required for the line conditions of your T1 line.
For example:
Buildout=0db
In most cases, you should not need to change the default value of 0db.

If you specify a value other than 0 db, the Pipeline 220 applies an attenuation to the T1 line, causing the line to lose power. (Repeaters boost the signal on a T1 line—if the Pipeline 220 is too close to a repeater, you might need to add some attenuation.)

- 7 Select Clock Source if this line should be used as the source for synchronous timing. In most cases, you should select Clock Source. Clearing the check box indicates to the Pipeline 220 that it should generate timing with its internal clock. Only in some back-to-back configurations should you configure the Pipeline 220 to generate timing.
- 8 Enter the number of channels assigned to this line by your carrier. You can enter any value from 1 to 24.

Enabling the settings

To enable your changes, you must:

- 1 Save the configuration by clicking the Save button on the lower-left side of the Configurator.
- 2 On the left side of the Configurator, click the System button.
- 3 On the right side of the Configurator, click the Info tab.
- 4 On the lower-right section of the Configurator, click the Reset button. The Configurator displays a dialog box asking for verification of the reset.
- 5 Select the Reset button and click OK. The Configurator displays a dialog box indicating the reset will occur in 60 seconds.
- 6 Click OK.

Configuring the serial WAN port

The Pipeline 220 has a built-in V.35 serial WAN DB-44 port. A serial WAN port provides a V.35/RS-449 WAN interface that is typically used to connect to a Frame Relay switch. The serial WAN data rate is determined by the clock speed received from the link. The maximum acceptable clock is 8 Mbps.

Signals to control the serial WAN data flow

The Activation parameter tells the Pipeline 220 which signals control the data flow through the serial WAN port. The DCE to which the serial WAN port is connected (such as a Frame Relay switch) determines how to set the value. Flow control is always handled by the Clear To Send (CTS) signal.

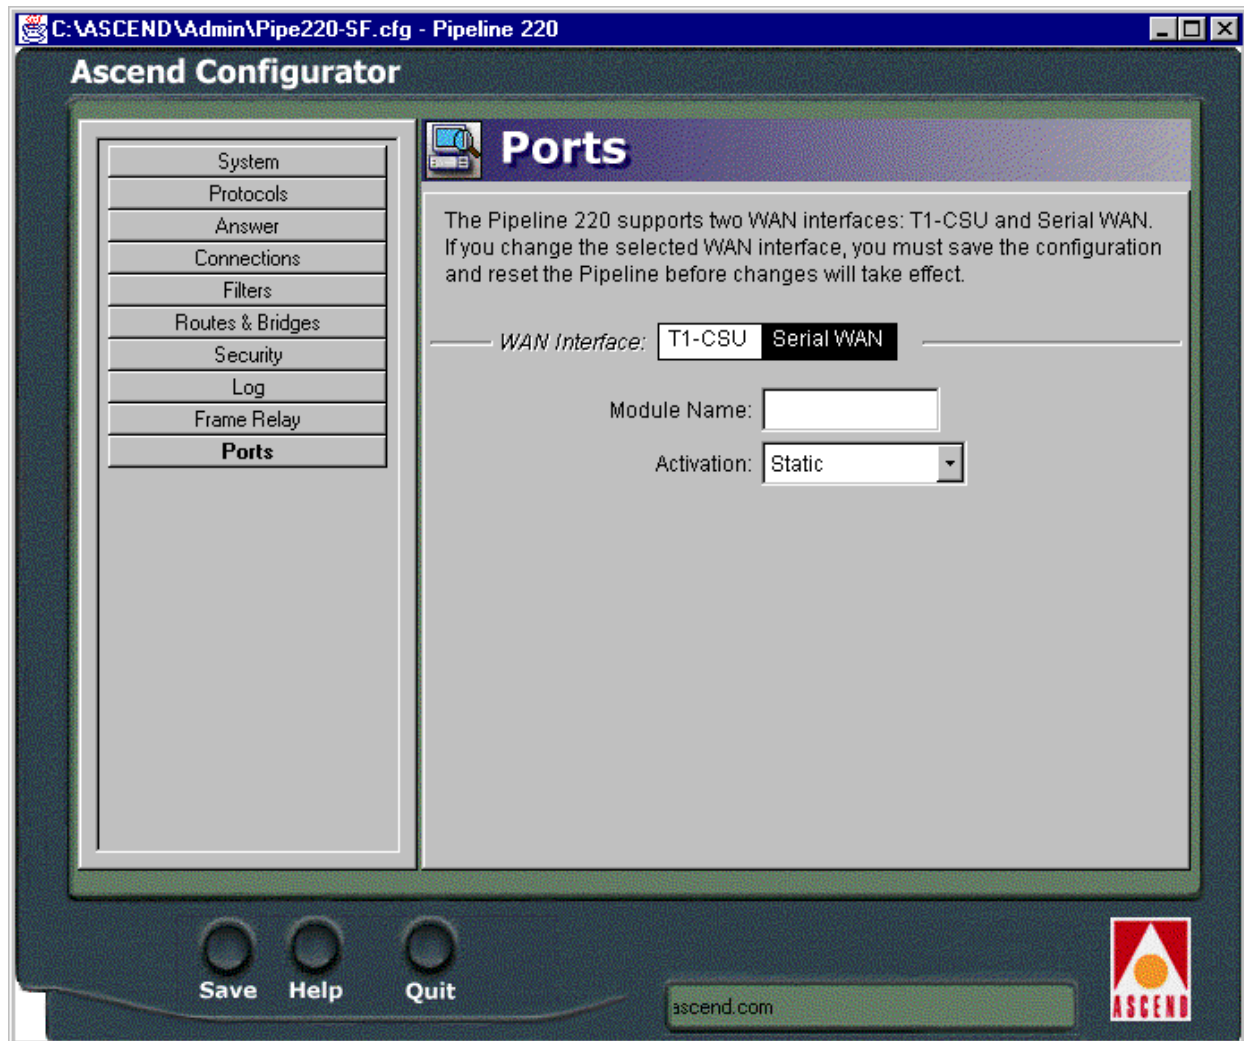
For details about each parameter, see the *Configurator Online Help*.

Configuring the serial WAN interface

Following is a procedure for configuring the serial WAN interface. This procedure configures the interface to connect to a Frame Relay switch that uses Static data flow:

- 1** On the left side of the Configurator, click the Ports button.
The right side of the Configurator now displays WAN Interface buttons.
- 2** Click the Serial WAN button.

The Configurator now displays parameters specific to the serial WAN interface:



- 3 Optionally, specify a name of up to 16 characters to Module Name.
For example:
Module Name=MainWANConn
- 4 Set Activation to Static.
- 5 Enable your changes, as described in “Enabling the settings” on page 4-4.

Configuring Frame Relay

This chapter covers the following topics:

Using the Pipeline 220 as a Frame Relay concentrator	5-1
Configuring the logical link to a Frame Relay switch	5-4
Configuring Connection profiles for Frame Relay	5-9

Using the Pipeline 220 as a Frame Relay concentrator

In a Frame Relay backbone, every access line connects directly to a Frame Relay switch. In the past, most connections to the Frame Relay network were relatively high speed, such as full T1 or E1 lines. But with recent changes in Frame Relay pricing, many sites now want to concentrate many low-speed dial-in connections into one high-speed nailed connection to a Frame Relay switch. When the Pipeline 220 is configured as a Frame Relay concentrator, it accepts incoming dial-in connections as usual, and forwards them out to a Frame Relay switch.




Figure 5-1. The Pipeline 220 operating as a Frame Relay concentrator

As a Frame Relay concentrator, the Pipeline 220 can accept up to 96 low-speed connections in North America or Japan, or 120 low-speed connections in Europe. If all of the Frame Relay connections are concentrated onto the single 2-Mbps serial WAN interface, the Pipeline 220 turns a single high-cost Frame Relay port on a traditional Frame Relay switch into approximately 100 operational ports.

Configuring the Pipeline 220 as a Frame Relay concentrator involves the following elements:

- An interface to the Frame Relay switch (usually nailed T1, nailed E1, or serial WAN)
- A logical datalink to the Frame Relay switch (defined in a Frame Relay profile)
- User connections (defined in Connection profiles)

For information about monitoring and managing Frame Relay, see “Monitoring Frame Relay connections” on page 16-32.

Kinds of physical network interfaces

The Pipeline 220 typically uses serial WAN, nailed T1, or nailed E1 to connect to a Frame Relay switch. For the details of configuring these interfaces, see Chapter 4, “Configuring the Pipeline 220 for WAN Access.”

Kinds of logical interfaces to a Frame Relay switch

The Pipeline 220 supports NNI, UNI-DCE, and UNI-DTE interfaces to the Frame Relay network.

Network to Network Interface (NNI)

With an NNI connection allows the Pipeline 220 appears to the switch to be a Frame Relay network interface. It performs both DTE and DCE link management, and allows two separate Frame Relay networks to connect via a common protocol. (To configure the interface, see “Configuring an NNI interface” on page 5-5.)




Figure 5-2. Network to Network interface (NNI) in a Pipeline 220 unit

User to Network Interface — Data Communications Equipment (UNI-DCE)

UNI is the interface between an end-user and a network end point (a router or a switch) on the Frame Relay network. In a UNI-DCE connection, the Pipeline 220 operates as a Frame Relay router communicating with a DTE device. To the DTE devices, it appears as a Frame Relay network end point. (To configure the interface, see “Configuring a UNI-DCE interface” on page 5-7.)




Figure 5-3. User to Network Interface-Data Communications Equipment (UNI-DCE)

User to Network Interface — Data Terminal Equipment (UNI-DTE)

In a UNI-DTE connection, the Pipeline 220 is configured as a UNI-DTE communicating with a Frame Relay switch. It acts as a Frame Relay feeder, and performs the DTE functions

specified for link management. To configure the interface, see “Configuring a UNI-DTE interface” on page 5-8.)




Figure 5-4. User to Network Interface - Data Terminal Equipment (UNI-DTE)

Types of Frame Relay connections

For Frame Relay connections, the Pipeline 220 supports gateway connections and Frame Relay circuits.

Gateway connections

With a gateway connection, the Pipeline 220 receives an incoming PPP call, examines the destination IP address, and brings up the appropriate Connection profile for that destination, as usual. If the Connection profile specifies Frame Relay encapsulation, the Frame Relay profile, and a DLCI, the Pipeline 220 encapsulates the packets in Frame Relay (RFC 1490), placing the DLCI in the headers, and forwards the data stream out to the Frame Relay switch. The Frame Relay switch uses the DLCI to route the frames. This is known as gateway mode.

Frame Relay circuits

A Frame Relay circuit is a permanent virtual circuit (PVC) segment that has two DLCI end points and a single Frame Relay profile. It requires two and only two DLCI numbers. Data is dropped if the circuit has only one DLCI. If more than two are defined, only two are used. A circuit is defined in two Connection profiles. Data coming in on the DLCI configured in the first Connection profile is switched to the DLCI configured in the second one.

Configuring the logical link to a Frame Relay switch

The Frame Relay profile specifies a link, usually across a single cable, to the Frame Relay network. This link can support many permanent virtual circuits (PVCs), each with a different endpoint.

Understanding the Frame Relay parameters

This section provides some background information about configuring the logical link to a Frame Relay switch. (For more detailed descriptions of the parameters, see the *Configurator Online Help*.)

Specifying a Frame Relay profile name and activating the profile

To provide access to the Frame Relay network, Connection profiles specify the name of the Frame Relay profile. Its name must be unique and cannot exceed 15 characters.

You must select the Frame Relay button > Active check box.

Bringing down the datalink when DLCIs are not active

If you select the Link Up check box, the data link comes up automatically and stays up even when the last DLCI has been removed. If this parameter is set to No, the datalink does not come up unless a Connection profile (specifying a DLCI) brings it up, and it shuts down after the last DLCI has been removed.

Note: You can start and drop Frame Relay datalink connections with the DO DIAL and DO HANGUP commands from the VT100 interface. DO DIAL brings up a datalink connection. DO HANGUP closes the link and any DLCIs on it. If LinkUp=Yes, DO HANGUP brings the link down, but it automatically restarts. A restart also occurs if there is a DLCI profile invoking the datalink. (For more information, see the system administration chapter in the *Configurator Online Help*.)

Defining the nailed connection to the switch

Nailed is the default value for Frame Relay connections. When the call type is Nailed, dial numbers and other telco options are N/A. You can specify switched if the Frame Relay switch allows dial-in. However, Frame Relay networks currently have no dial-out connection capability. The two types of data service available are 64K and 56K.

Specifying the type of Frame Relay interface

You can set the FR Type parameter to NNI (for an NNI interface to the switch), DCE (for a UNI-DCE interface), or DTE (for a UNI-DTE interface). (For a description of the interfaces, see “Kinds of logical interfaces to a Frame Relay switch” on page 5-2.)

Link management protocol

The Link Mgmt setting may be None (no link management), T1.617D (for T1.617 Annex D), or Q.933A (for Q.933 Annex A).

Frame Relay timers and event counts

To locate Frame Relay timers and event counts, you click the Frame Relay button, then the Link Management button, to display the following fields:

- “Request Full Status Report” (known in RFC 1490 as N391) specifies the interval at which the Pipeline 220 requests a Full Status Report from the frame relay switch (between 1 and 255 seconds). It does not apply if FR Type is DCE.
- “Error Threshold” (known in RFC 1490 as N392) specifies the number of errors, during the number of monitored events specified in the “Event count” field, that causes the network-side to declare the frame relay connection inactive. The value in “Error Threshold” should be less than that of “Event Count” (which may be from 1 to 10). The fields for DCE do not apply when FR Type (displayed by clicking the General tab) is DTE. The fields for DTE do not apply when FR Type is DCE.
- “Event count” (known in RFC 1490 as N393) specifies the monitored event count (from 1 to 10). The fields for DCE do not apply when FR Type (displayed by clicking the General tab) is DTE. The fields for DTE do not apply when FR Type is DCE.
- “Delay Between Status Inquiry Messages” (known in RFC 1490 as T391) specifies the Link Integrity Verification polling timer (which may be from 5 to 30 seconds). Its value should be less than that of the “Delay to wait for messages before recording an error” field. It does not apply when FR Type is DCE.
- “Delay to wait for messages before recording an error” (known in RFC 1490 as T392) specifies the maximum time allowed between Status Enquiry messages (from 5 to 30 seconds). An error is recorded if no Status Enquiry is received within the value you specify. This parameter does not apply when FR Type is DTE.

Maximum Receive Units (MRU)

The MRU parameter specifies the maximum number of bytes the Pipeline 220 can receive in a single packet across the link. Usually the default of 1532 is the correct setting. However, the far-end device might require a lower number.

Examples of Frame Relay profile configurations

This section shows a sample Frame Relay profile configuration for each type of Frame Relay interface (NNI, UNI-DCE, and UNI-DTE).

Configuring an NNI interface

In the following example, the Pipeline 220 has a nailed connection to another Frame Relay switch and will be configured with an NNI interface to that switch. Figure 5-5 shows the connection.




Figure 5-5. NNI interface to another switch

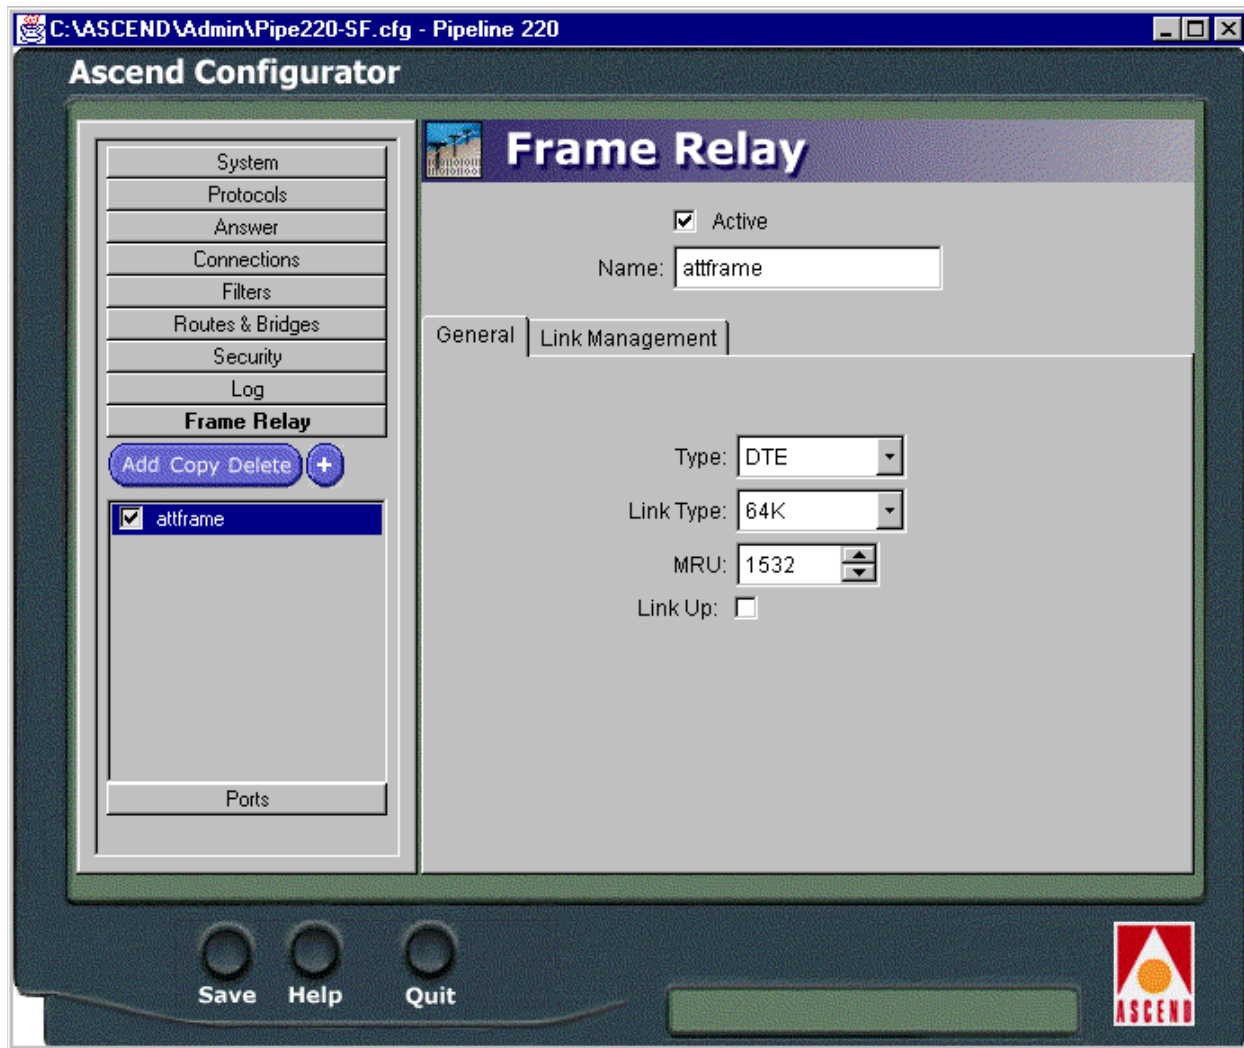
Configuring Frame Relay

Configuring the logical link to a Frame Relay switch

To configure the Frame Relay profile for this NNI interface:

- 1 On the left side of the Configurator, click the Frame Relay button.
- 2 On the right side of the Configurator, click the General button.

System-wide Frame Relay parameters appear on the right side of the Configurator:



- 3 Assign the profile a name and select the Active check box.
- 4 Set Type to NNI.
- 5 Set Link Type as directed by your Frame Relay provider.
- 6 On the left side of the parameter, click the Link Management button.
- 7 Specify the link management protocol and its configuration parameters, as directed by your Frame Relay provider.
- 8 Enable the new settings, as described in “Saving the settings” on page 5-6.


Saving the settings

The Pipeline 220 settings you configure with the Ascend Configurator do not take effect until you upload them to the Pipeline 220. You use the Save command to upload the changes.

Alternatively, you can save the settings to a text file for subsequent use, or to another device. To save the settings:

- 1 In the lower-left corner of the Configurator, click Save.

A dialog box appears, prompting you for a save method:



- 2 Select one of the following:
 - If you are uploading the configuration to the same device, select “Upload changes to <Pipeline 220 name>.”
 - If you are saving the configuration to a text file on your computer, select “Save a copy under a new filename” and enter the filename to which the configuration should be saved.
 - If you are uploading the configuration to a different device, select “Upload to the Ascend product at:” and enter the IP address of the device.
- 3 Click Save.

If you selected “Upload changes to <Pipeline 220 name>” in step 2, the Ascend

Configurator uploads the configuration file to the Pipeline 220 via TFTP. If you selected one of the other options, the Configurator sends the file to the location you specified. If you specified a remote device, transmission is via TFTP.

Configuring a UNI-DCE interface

In the following example, the Pipeline 220 has a nailed connection to customer premises equipment (CPE) and will be configured with a UNI-DCE interface to that equipment. Figure 5-6 shows the connection.




Figure 5-6. UNI-DCE interface to an end-point (DTE)

To configure the Frame Relay profile for this UNI-DCE interface:

- 1 On the left side of the Configurator, click the Frame Relay button.
- 2 On the right side of the Configurator, click the General button.
System-wide Frame Relay parameters appear on the right side of the Configurator.
- 3 Assign the profile a name and select the Active check box.
- 4 Set Type to DCE.
- 5 Set Link Type as directed by your Frame Relay provider.
- 6 On the left side of the parameter, click the Link Management button
- 7 Specify the link management protocol and its configuration parameters as directed by your Frame Relay provider.
- 8 Enable the new settings, as described in “Saving the settings” on page 5-6.

Configuring a UNI-DTE interface

In this example, the Pipeline 220 has a nailed connection to a Frame Relay switch configured as a DCE, and will be configured with a UNI-DTE interface to that switch. Figure 5-7 shows the connection.




Figure 5-7. UNI-DTE interface to a Frame Relay switch

To configure the Frame Relay profile for this UNI-DTE link:

- 1 On the left side of the Configurator, click the Frame Relay button.
- 2 On the right side of the Configurator, click the General button.
System-wide Frame Relay parameters appear on the right side of the Configurator.
- 3 Assign the profile a name and select the Active check box.
- 4 Set Type to DTE.
- 5 Set Link Type as directed by your Frame Relay provider.
- 6 On the left side of the parameter, click the Link Management button.
- 7 Specify the link management protocol and its configuration parameters as directed by your Frame Relay provider.
- 8 Enable the new settings, as described in “Saving the settings” on page 5-6.

Configuring Connection profiles for Frame Relay

For each connection that uses the Frame Relay link, you must configure a Connection profile that specifies the Frame Relay profile name as the datalink between the Pipeline 220 and the Frame Relay network. The Frame Relay profile specifies the physical datalink connection. Each Connection profile specifies each PVC. Every Connection profile specifies a unique DLCI, while specifying the identical Frame Relay profile.

Understanding the Frame Relay connection parameters

This section provides some background information about configuring a Connection profile that is used in combination with the Frame Relay connection. (For more detailed descriptions of the parameters, see *Configurator Online Help*.)

Gateway connections

Gateway connections require FR encapsulation, a Frame Relay profile name, and a DLCI. Your Frame Relay provider gives you the DLCI to assign to each connection.

The far end specified in a Frame Relay-encapsulated Connection profile lies at the end of a PVC, whose first hop is known by the DLCI named in the Connection profile. The Pipeline 220 does not allow you to enter duplicate DLCIs, except when they are carried by separate physical links specified in different Frame Relay profiles.

Frame Relay circuits

A circuit is a PVC segment configured in two Connection profiles. Data coming in on the DLCI configured in one Connection profile is switched to the DLCI configured in the other. Data is dropped if the circuit has only one DLCI. If more than two Connection profiles specify the same circuit name, only two of them are used.

In a circuit, both Connection profiles must specify FR_CIR encapsulation, and both must specify the same circuit name. Each Connection profile must specify a unique DLCI. The Pipeline 220 does not allow you to enter duplicate DLCIs for a single profile.

Examples of connection configurations

This section shows sample Connection profile configurations for Frame Relay gateway and circuit configurations.

Configuring a Frame Relay gateway connection

This example shows you how to configure a Frame Relay gateway connection. It assumes that dial-in users who need to reach the distant IP network have valid Connection profiles. The sample procedure is for the Connection profile for the connection to the destination. The profile assigns a DLCI and passes the data stream out to a Frame Relay switch. Figure 5-8 shows the network.





Figure 5-8. Gateway connections

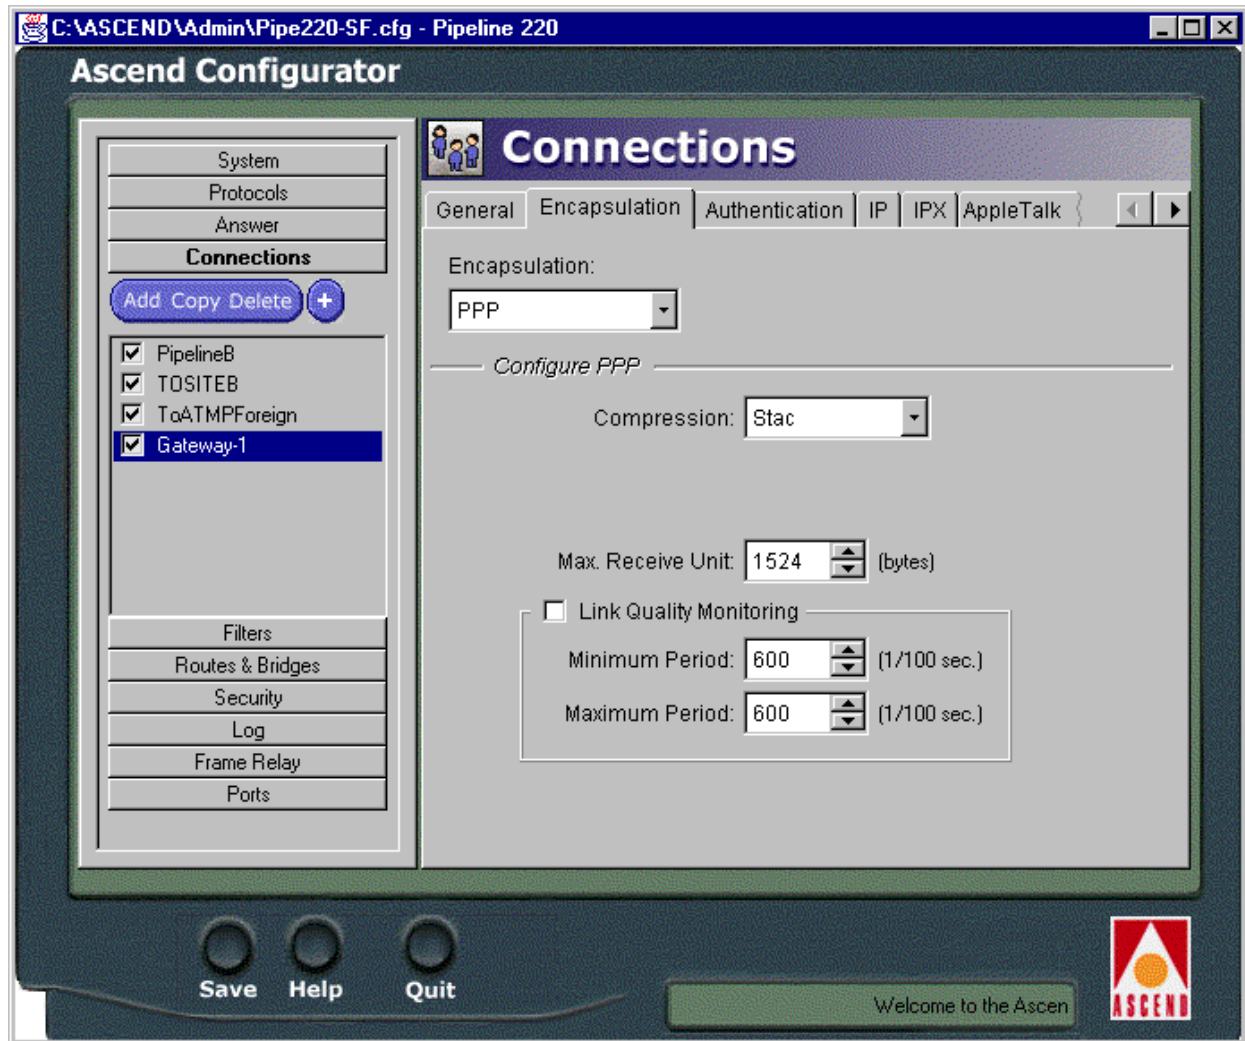
In this example, the Pipeline 220 uses a Frame Relay profile named ATT-NNI for communication with a remote Frame Relay switch. To configure the corresponding Connection profile:

- 1 On the left side of the Configurator, click the Connections button.
The Add/Copy/Delete window appears on the lower-left of the Configurator.
- 2 In the Add/Copy/Delete window, click Add.
The Configurator displays a dialog box prompting you for a name of the new connection:



- 3 Enter a name for the connection, then click OK. For example:
Name=Gateway-1

The Connection profile appears on the right side of the Configurator, displaying the name and all other values at their factory default settings:



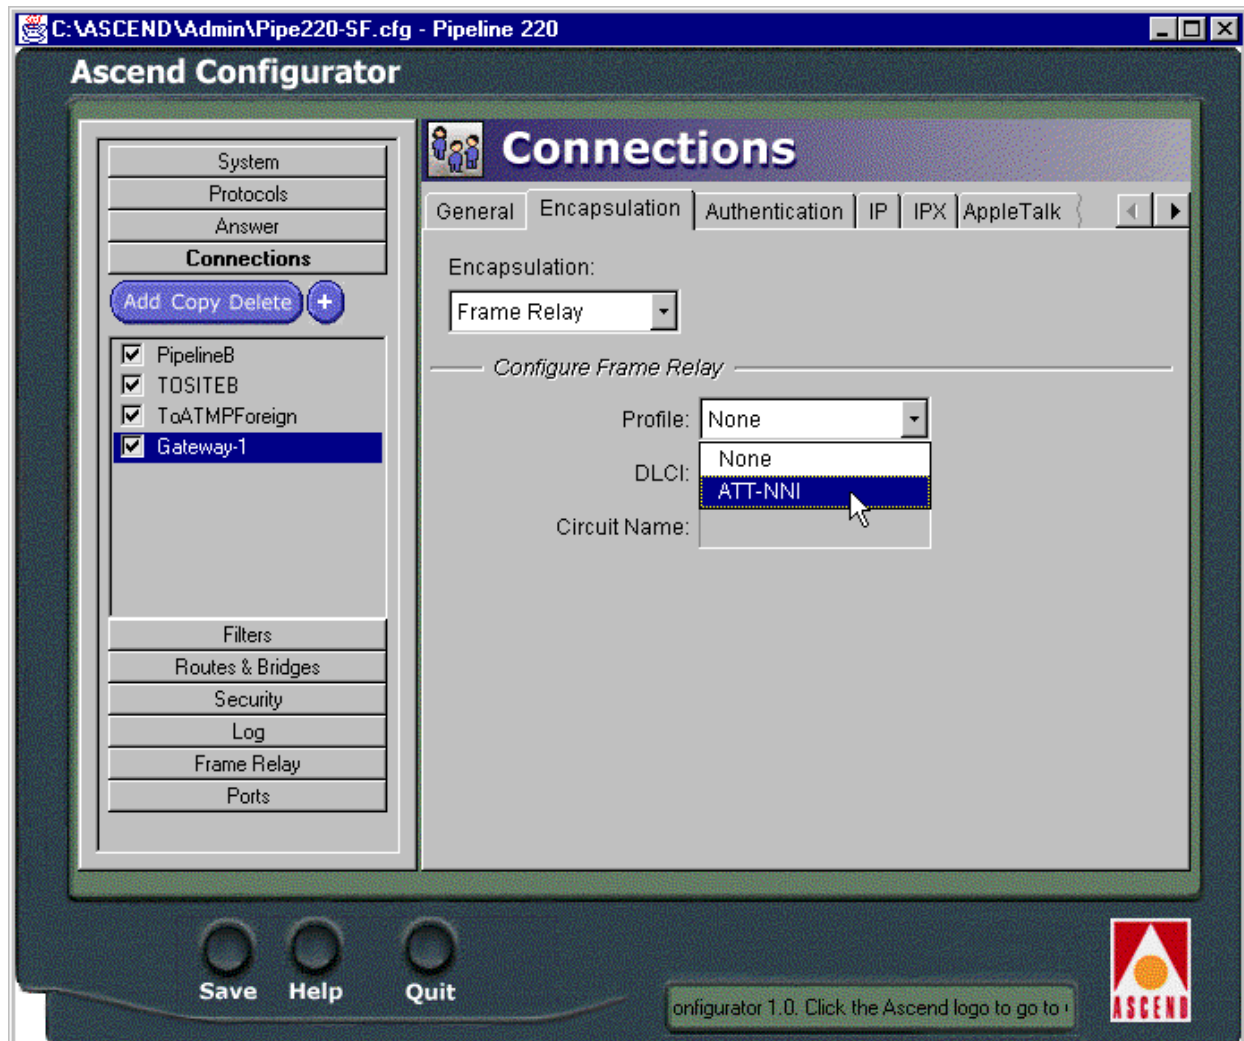
- 4 Click the IP tab to configure IP routing options.
- 5 Select the Enable IP Routing check box.
- 6 Click the Addresses IP-options button.
- 7 Set the IP Address parameter to the address the remote device's Ethernet interface:
IP Address=10.9.8.10
Subnet Mask=22
- 8 Click the Encapsulation tab.

Configuring Frame Relay

Configuring Connection profiles for Frame Relay

- 9 Set Encapsulation to Frame Relay.

Frame Relay parameters appear in the lower-right section of the Configurator:



- 10 For the Profile parameter, select ATT-NNI.
- 11 Set DLCI to the DLCI number supplied by your Frame Relay provider.
- 12 Enable the new settings, as described in “Saving the settings” on page 5-6.

Configuring a Frame Relay circuit

The procedure in this example configures a circuit between UNI-DCE and NNI datalinks. A circuit between any two interfaces within the Pipeline 220 would be configured in much the same way. Figure 5-9 shows the sample network.





Figure 5-9. A Frame Relay circuit

The Frame Relay profile for the UNI-DCE interface in the Pipeline 220 is named ATT-DCE. For the NNI interface, the Frame Relay profile is named ATT-NNI.

Configuring the Connection profile for the UNI-DCE interface

To configure the first Connection profile for this circuit:

- 1 On the left side of the Configurator, click the Connections button.
The Add/Copy/Delete window appears in the lower-left side of the Configurator.
- 2 In the Add/Copy/Delete window, click Add.
The Configurator displays a dialog box prompting you for a name of the new connection:



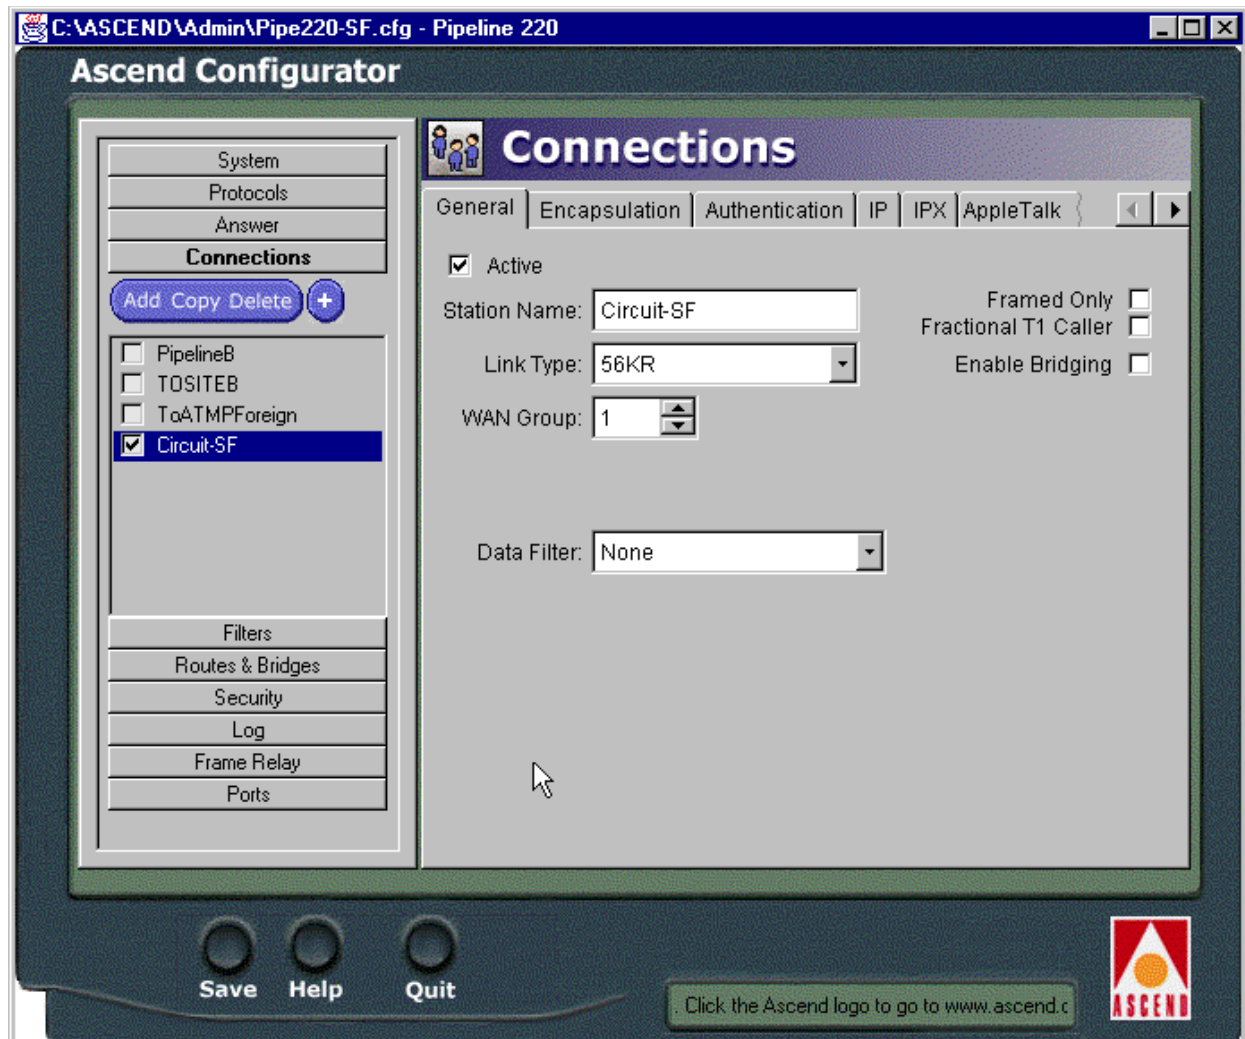
Configuring Frame Relay

Configuring Connection profiles for Frame Relay

- 3 Enter a name for the connection, then click OK. For example:

Name=Circuit-SF

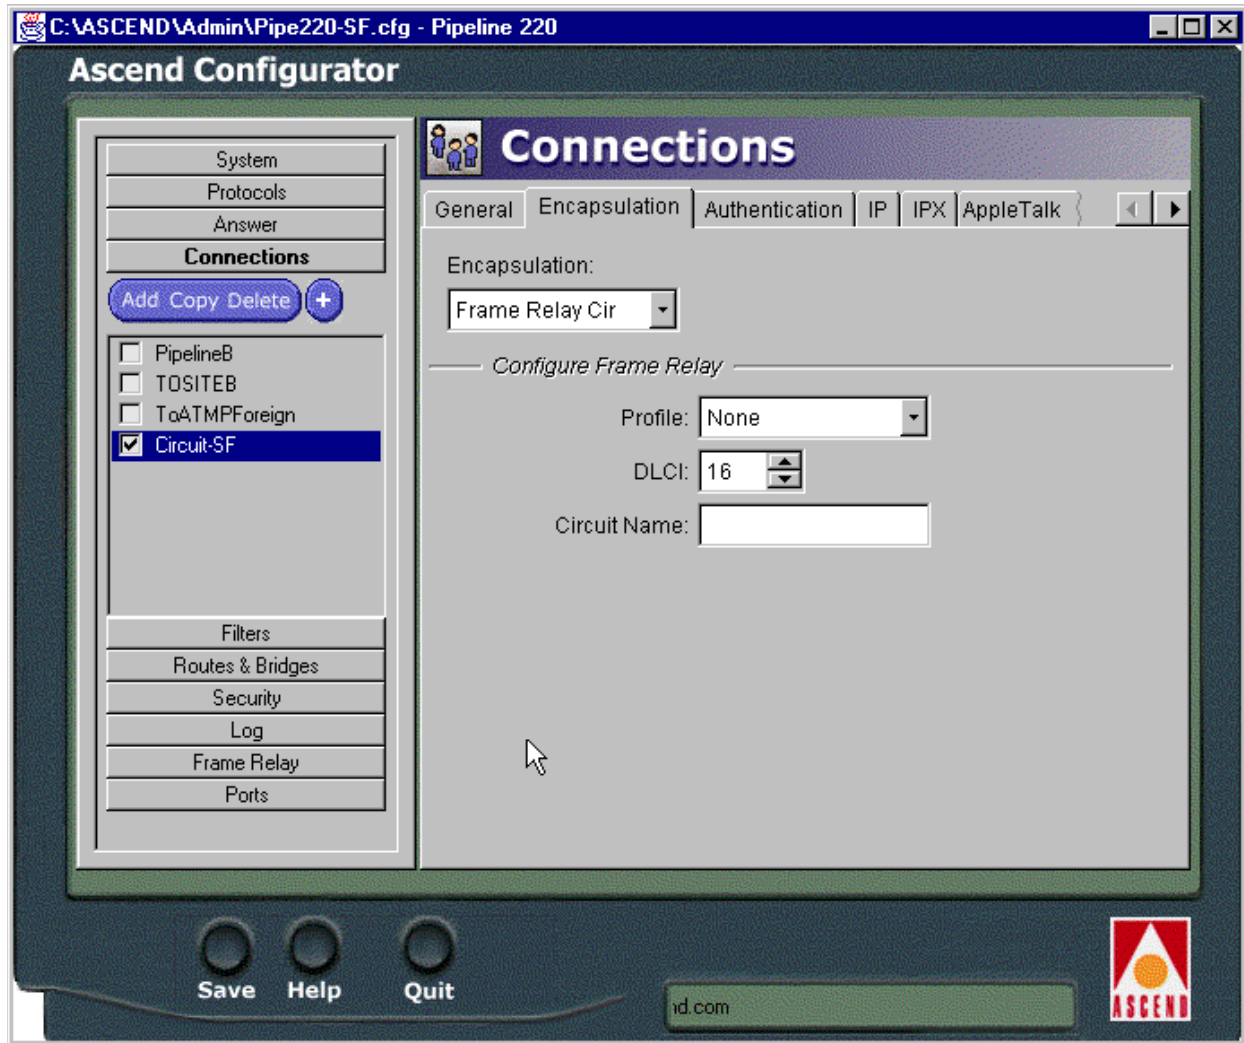
The Connection profile appears on the right side of the Configurator, displaying the name and all other values at their factory default settings:



- 4 Click the Encapsulation tab.

- 5 Set Encapsulation to Frame Relay Cir.

Frame Relay Circuit parameters appear in the lower-right section of the Configurator:



- 6 For the Profile parameter, select ATT-DCE.
- 7 Set DLCI to the DLCI number supplied by your Frame Relay provider.
- 8 Set Circuit to a descriptive name for the Frame Relay circuit.
This name must be the same in both Connection profiles you create for this circuit.

Configuring the Connection profile for the NNI interface


To configure the first Connection profile for this circuit:

- 1 On the left side of the Configurator, click the Connections button.
The Add/Copy/Delete window appears in the lower-left side of the Configurator.
- 2 In the Add/Copy/Delete window, click Add.

Configuring Frame Relay

Configuring Connection profiles for Frame Relay

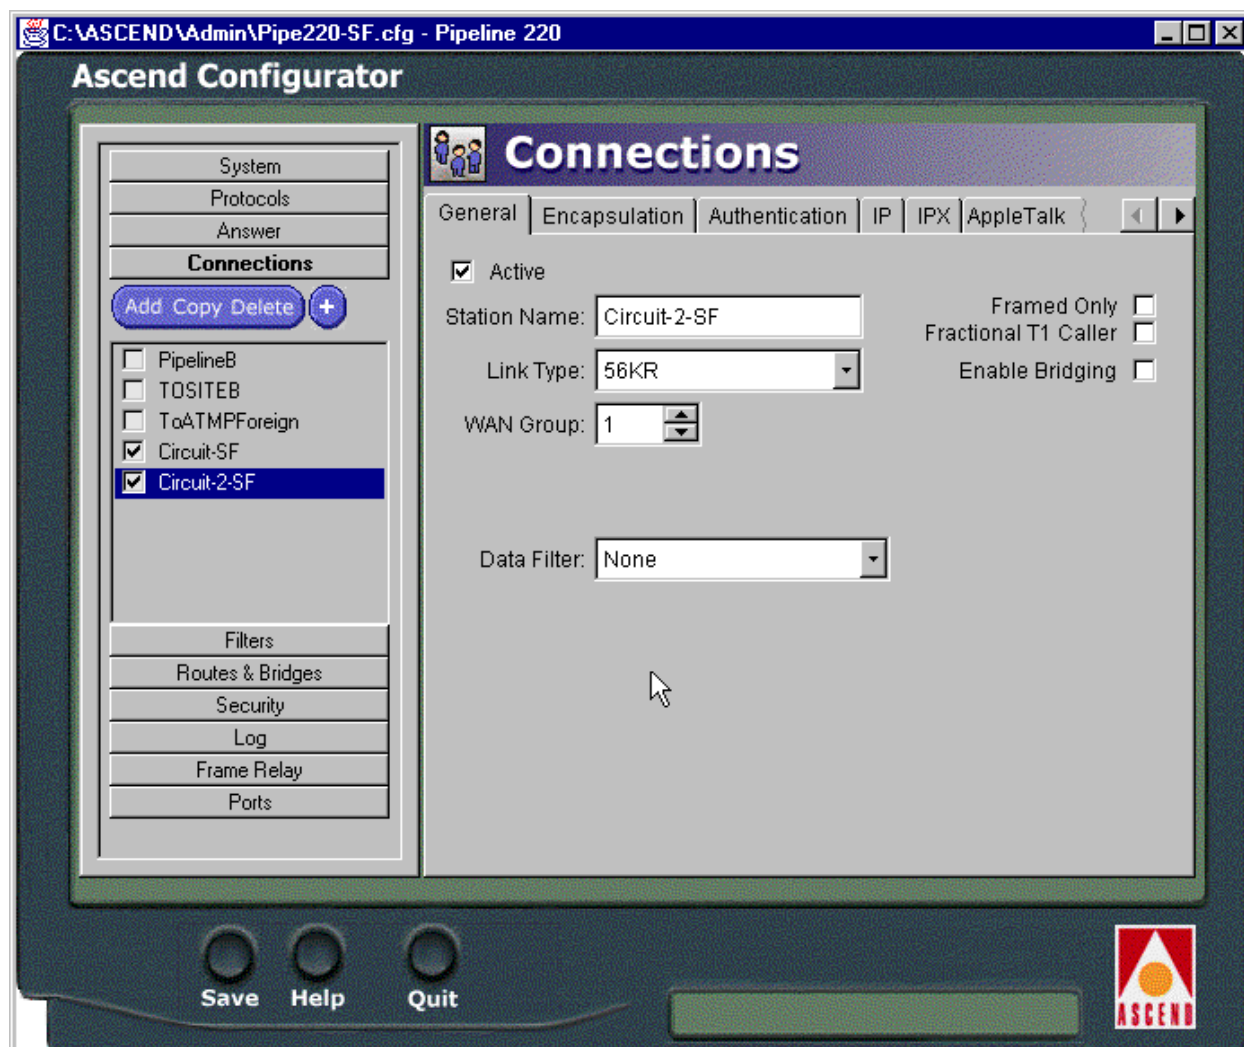
The Configurator displays a dialog box prompting you for a name of the new connection:



- 3 Enter a name for the connection, then click OK. For example:

Name=Circuit-SF

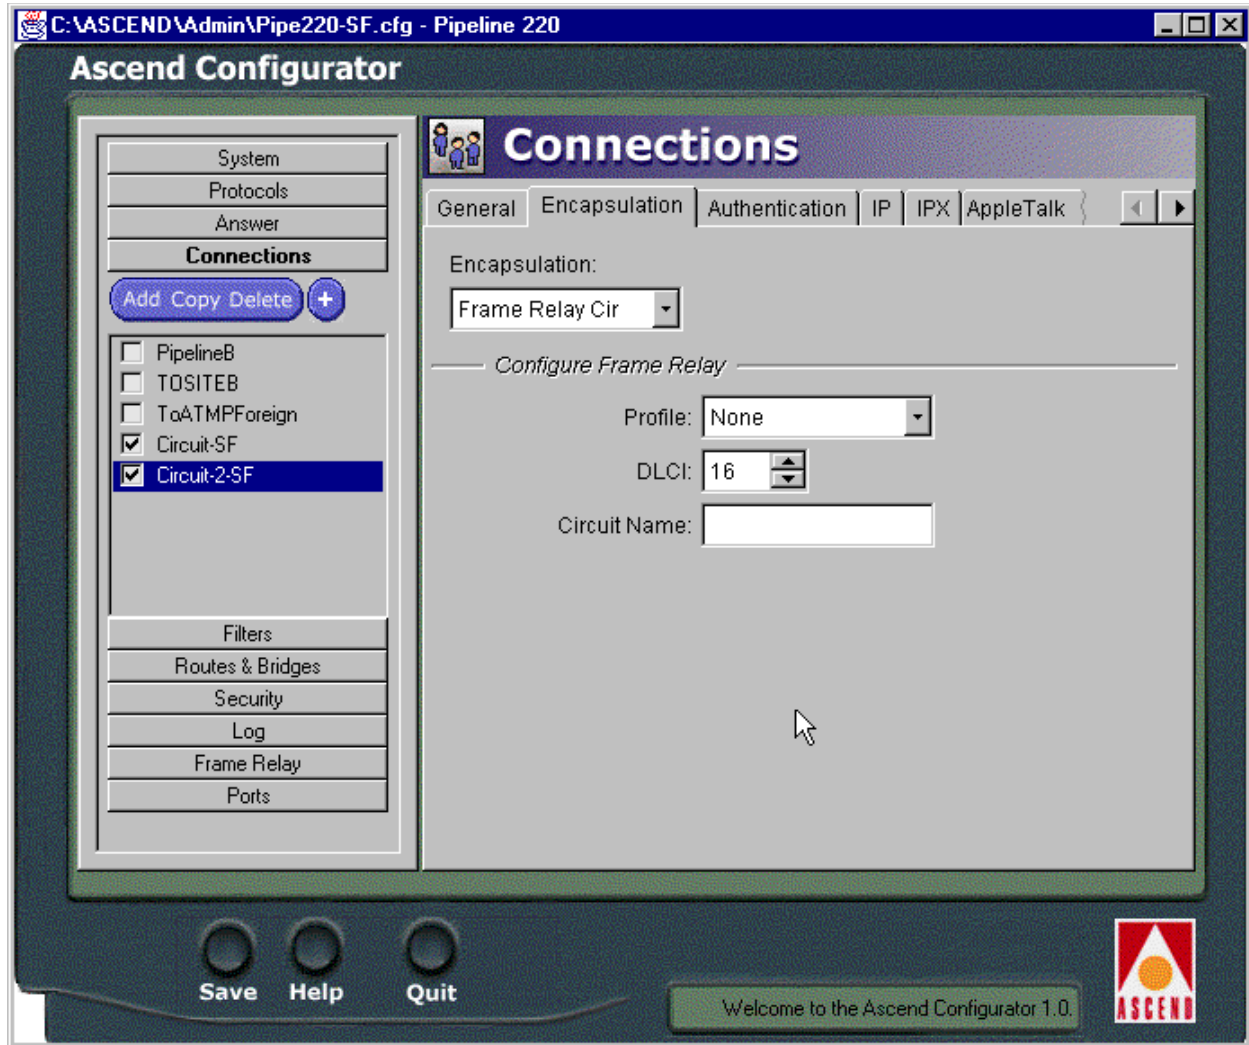
The Connection profile appears on the right side of the Configurator, displaying the name and all other values at their factory default settings:



- 4 Click the Encapsulation tab.

5 Set Encapsulation to Frame Relay Cir.

Frame Relay Circuit parameters appear in the lower-right section of the Configurator:



- 6 For the Profile parameter, select ATT-NNI.
- 7 Set DLCI to the DLCI number your Frame Relay provider gives you.
- 8 Set Circuit to a descriptive name for the Frame Relay circuit.
This name must match for both Connection profiles you create for this circuit.
- 9 Save the new settings, as described in "Saving the settings" on page 5-6.

Configuring IP Routing

This chapter covers the following topics:

Introduction to IP routing and interfaces	6-1
Configuring the local IP network setup	6-8
Configuring IP routing connections.	6-17
Configuring IP routes and preferences.	6-27
Configuring the Pipeline 220 for dynamic route updates	6-34
Syslog services.	6-37

Introduction to IP routing and interfaces

The first task described in this chapter, setting up the IP network, involves setting parameters in the Pipeline 220 unit's Ethernet profile. The parameters define the unit's Ethernet IP interface, network services (such as DNS), and routing policies.

In the next task, configuring IP routing connections, you configure Connection profiles (or similar profiles in an external authentication server) to define destinations across WAN interfaces and add routes to the routing table.

For configuring IP routes and preferences and configuring the Pipeline 220 for dynamic route updates, you configure the IP profile and individual Connection profiles to set up the IP routing table, which determines the paths over which IP packets are forwarded and specifies the connections to be brought up.

To perform the tasks described in this chapter, you have to understand how the Pipeline 220 uses IP addresses and subnet masks, IP routes, and IP interfaces.

IP addresses and subnet masks

In the Pipeline 220, you specify IP addresses in dotted decimal format (not hexadecimal). If you specify no subnet mask, the Pipeline 220 assumes a default mask on the basis of address

class. The default subnet mask is the default number of network bits for the address's class. Table 6-1 shows the classes and the default number of network bits for each class.

Table 6-1. IP address classes and default subnet masks

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

For example, a class C address such as 198.5.248.40 has 24 network bits, so its default mask is 24. The 24 network bits leave 8 bits for the host portion of the address. So one class C network can support up to 253 hosts.



Figure 6-1. A class C IP address

For specifying a different subnet mask, the Pipeline 220 supports a modifier that specifies the total number of network bits in the address. For example:

```
IP address = 198.5.248.40
Mask = 255.255.255.248
```

In the example address shown above, the mask specification indicates that 29 bits of the address will be used to specify the network. This is commonly referred to as a 29-bit subnet. The three remaining bits specify unique hosts.

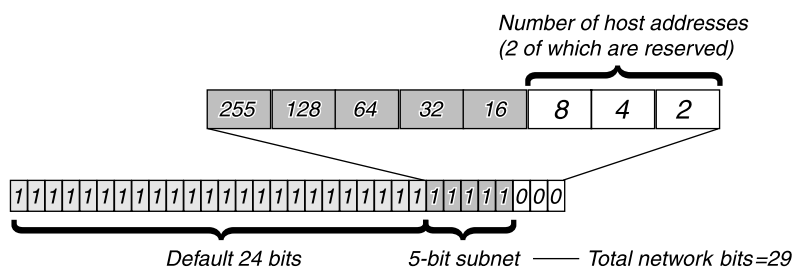


Figure 6-2. A 29-bit subnet mask and number of supported hosts

Three available bits allow eight possible bit combinations. Of the eight possible host addresses, two are reserved, as follows:

```
000 — Reserved for the network (base address)
001
010
100
110
101
```


011

111 — Reserved for the broadcast address of the subnet

Zero subnets

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 and OSPF treat these so-called zero subnetworks the same as any other network. You should decide whether or not to support and configure zero subnetworks for your environment. If you configure them in some cases and treat them as unsupported in other cases, you will encounter routing problems.

Table 6-2 shows how the standard subnet address format relates to Ascend notation for a class C network number.

Table 6-2. Standard subnet masks

Subnet mask	Number of host addresses
255.255.255.0	254 hosts + 1 broadcast, 1 network (base)
255.255.255.128	126 hosts + 1 broadcast, 1 network (base)
255.255.255.192	62 hosts + 1 broadcast, 1 network (base)
255.255.255.224	30 hosts + 1 broadcast, 1 network (base)
255.255.255.240	14 hosts + 1 broadcast, 1 network (base)
255.255.255.248	6 hosts + 1 broadcast, 1 network (base)
255.255.255.252	2 hosts + 1 broadcast, 1 network (base)
255.255.255.254	invalid netmask (no hosts)
255.255.255.255	1 host — a host route

The broadcast address of any subnet has the host portion of the IP address set to all ones. The network address (or base address) represents the network itself, with the host portion of the IP address set to all zeros. Therefore, these two addresses define the address range of the subnet. For example, if the Pipeline 220 configuration assigns the following address to a remote router:

IP address = 198.5.248.120

Mask = 255.255.255.248

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

198.5.248.120 — 198.5.248.127

A host route is a special case IP address with a subnet mask of 32 bits. It has a subnet mask of 255 . 255 . 255 . 255.

IP routes

At system startup, the Pipeline 220 builds an IP routing table that contains configured routes. When the system is up, it can use routing protocols such as RIP or OSPF to learn additional routes dynamically.

In each routing table entry, the Destination field specifies a destination network address that may appear in IP packets, and the Gateway field specifies the address of the next-hop router to reach that destination.

How the Pipeline 220 uses the routing table

The Pipeline 220 relies on the routing table to forward IP packets, as follows:

- If the Pipeline 220 finds a routing table entry whose Destination field matches the destination address in a packet, it routes the packet to the specified next-hop router, whether through its WAN interface or through its Ethernet interface.
- If the Pipeline 220 does not find a matching entry, it looks for the Default route, which is identified in the routing table by a destination of 0.0.0.0. If that route has a specified next-hop router, it forwards the packet to that router.
- If the Pipeline 220 does not find a matching entry or does not have a valid Default route, it drops the packet.

Static and dynamic routes

A static route is a manually configured path from one network to another, which specifies the destination network and the gateway (router) to use to get to that network.

- Each Routes profile specifies one static route. If a path to a destination must be reliable, the administrator often configures more than one path (that is, specifies one or more secondary routes), in which case the Pipeline 220 chooses the route on the basis of assigned metrics and availability.
- The Protocols button > IP tab > Addresses button specifies a static connected route, which states “to reach system-A, send packets out this interface”.
- Each IP-routing Connection profile specifies a static route that states “to reach system-A, send packets out this interface to system-B”, where system-B is another router.

A dynamic route is a path, to another network, that is learned from another IP router rather than configured in one of the Pipeline 220's local profiles. Routers that use RIP broadcast their entire routing table every 30 seconds, updating other routers about the usability of particular routes. Hosts that run ICMP can also send ICMP Redirects to offer a better path to a destination network. OSPF routers propagate link-state changes as they occur. Routing protocols such as RIP and OSPF all use some mechanism to propagate routing information and changes through the routing environment.

Route preferences and metrics

The Pipeline 220 supports route preferences, because different protocols have different criteria for assigning route metrics. For example, RIP is a distance-vector protocol, which uses a virtual hop count to select the shortest route to a destination network. OSPF is a link-state protocol, which means that OSPF can take into account a variety of link conditions, such as the reliability or speed of the link, when determining the best path to a destination network.

When choosing a route to put into the routing table, the router first compares preference values, preferring the lowest number. If the preference values are equal, the router compares the metric fields and uses the route with the lowest metric. Following are the preference values for the various types of routes:

- Connected routes have a default preference of 0.
- OSPF routes have a default preference of 10.
- ICMP redirects have a default preference of 30.
- RIP routes have a default preference of 100.
- Static routes have a default preference of 100.
- ATMP, PPTP routes have a default preference of 100.

Pipeline 220 IP interfaces

The Pipeline 220 must have at least one system-based IP interface (on Ethernet) to support IP routing. It also creates several internal interfaces at system startup.

At system startup, the Pipeline 220 creates its Ethernet and internal IP interfaces.

Ethernet interfaces

The following example displays the routing table for a Pipeline 220 configured to enable IP routing:

```
** Ascend Pipeline Terminal Server **
```

```
ascend% iproute show
```

Destination	Gateway	IF	Flg	Pref	Met	Use	Age
10.10.0.0/16	-	ie0	C	0	0	3	222
10.10.10.2/32	-	local	CP	0	0	0	222
127.0.0.0/8	-	bh0	CP	0	0	0	222
127.0.0.1/32	-	local	CP	0	0	0	222
127.0.0.2/32	-	rj0	CP	0	0	0	222
224.0.0.0/4	-	mcast	CP	0	0	0	222
224.0.0.1/32	-	local	CP	0	0	0	222
224.0.0.2/32	-	local	CP	0	0	0	222
224.0.0.5/32	-	local	CP	0	0	0	222
224.0.0.6/32	-	local	CP	0	0	0	222
224.0.0.9/32	-	local	CP	0	0	0	222
255.255.255.255/32	-	ie0	CP	0	0	0	222

The Ethernet interface has the IP address 10.10.10.2 (with a subnet mask of 255.255.0.0). No Connection profiles or static routes are configured.

Following are descriptions of the interfaces created at startup:

- The Ethernet IP interface, labeled `ie0`, is always active, because it is always connected. Its IP address is assigned in Protocols button > IP tab > Addresses button > IP Address.

The Pipeline 220 creates two routing table entries: one with a destination of the network (labeled `ie0`), and the other with a destination of the Pipeline 220 (labeled `local`).

Note: Remember that the Pipeline 220 has two separate Ethernet ports. Each of those ports support two separate IP addresses.

- The black-hole (`bh0`) interface is always up. The black-hole address is 127.0.0.3. Packets routed to this interface are discarded silently.
- The loopback (labeled `local`) interface is always up. The loopback address is 127.0.0.1/32.
- The reject (labeled `rj0`) interface is always up. The reject address is 127.0.0.2. Packets routed to this interface are sent back to the source address with an ICMP “host unreachable” message.
- Multicast interfaces have a destination address with a value of 224 for the first octet. For information about multicast addresses, see Chapter 9, “Setting Up IP Multicast Forwarding.”
- Not shown in the example is an inactive interface. It is created when you configure a Connection profile. The inactive interface is where all routes point when their WAN connections are down. The inactive interface label is `wanidle0`.

WAN IP interfaces

WAN interfaces are created as they are brought up. WAN interfaces are labeled `wanN`, where *N* is a number assigned in the order in which the interfaces become active. The WAN IP address can be a local address assigned dynamically when the caller logs in, an address on a subnet of the local network, or a unique IP network address for a remote device.

Numbered interfaces

The Pipeline 220 can operate as both a system-based and interface-based router. Interface-based routing uses numbered interfaces. Some routers or applications require numbered interfaces, and some sites use them for trouble-shooting 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 220 to operate in much the same way as a multihomed Internet host.

Figure 6-3 shows a sample interface-based routing connection.




Figure 6-3. Interface-based routing example

The IP addresses 10.5.6.7 and 10.5.6.8 are assigned to the WAN interfaces. The site A Pipeline 220 routes packets to the remote network 10.7.8.0 by means of the addresses assigned the WAN interfaces.

With system-based routing, these addresses are not assigned. The site A Pipeline 220 routes packets to the remote network on the basis of the WAN interface it created when the connection was brought up, rather than a configured IP address.

Interface-based routing means that in addition to the system-wide IP configuration, the Pipeline 220 and the far end of the link have link-specific IP addresses, for which you specify the following parameters:

- Connections button > IP tab > Addresses button > Interface Address (the link-specific address for the Pipeline 220)
- Connections button > IP tab > Addresses button > WAN Alias (the far end link-specific address)

Or, you may omit the remote side's system-based IP address from the Connection profile and use interface-based routing exclusively. This is an appropriate mechanism, for example, if the remote system is on a backbone net that 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 link-specific IP addresses are specified in the following parameters:

- Connections button > IP tab > Addresses button > Interface Address (the near end numbered interface)
- Connections button > IP tab > Addresses button > IP Address (the far end numbered interface)

Note that IP Address must always be filled in, so if the only known address is the interface address, you must place it in the IP Address parameter rather than the WAN Alias parameter. In this case, a host route is created to the interface address (IP address), a net route is created to the subnet of the remote interface, and incoming calls must report their IP addresses as the value in the IP Address parameter.

It is also possible, although not recommended, to specify the local numbered interface (Interface Address) and use the far end device's system-wide IP address (IP Address). In this case, the remote interface must have an address on the same subnet as the local, numbered interface.

If a Pipeline 220 is using a numbered interface, note the following differences and similarities in operation, compared to unnumbered (system-based) routing:

- IP packets generated in the Pipeline 220 and sent to the remote address will have an IP source address corresponding to the numbered interface, not the system-wide (Ethernet) address.
- The Pipeline 220 adds all numbered interfaces to its routing table as host routes.
- The Pipeline 220 accepts IP packets addressed to a numbered interface, considering them to be destined for the Pipeline 220 itself. (The packet may actually arrive over any interface, and the numbered interface corresponding to the packet's destination address need not be active.)

Configuring the local IP network setup

The IP tab displayed when you select the Protocols button configures system parameters that affect all IP interfaces in the Pipeline 220.

Understanding the IP network parameters

This section provides some background information about the IP network configuration. The information is organized by functionality rather than by parameter.

Primary IP address for each Ethernet interface

The IP Address parameter specifies the Pipeline 220 unit's IP address for each local Ethernet interface. When specifying IP addresses for the Pipeline 220's Ethernet interfaces, you must specify the subnet mask. IP address and subnet mask are required settings for the Pipeline 220 to operate as an IP router.

Second IP address for each Ethernet interface

The Pipeline 220 can assign two unique IP addresses to *each* physical Ethernet port and route between them. This feature, referred to as *dual IP*, can give the Pipeline 220 a logical interface on two networks or subnets on the same backbone.

Usually, devices connected to the same physical wire all belong to the same IP network. With dual IP, a single wire can support two separate IP networks, with devices on the wire assigned to one network or the other and communicating by routing through the Pipeline 220.

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. When the routers have a direct connection to the subnet as well as to the backbone network, they route packets to the subnet and include 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 placeholder while you are making the transition in other network equipment.

Figure 6-4 shows an example IP network to which a Pipeline 220 is connected:




Figure 6-4. Sample dual IP network

Two IP addresses are assigned to each of the Pipeline 220's Ethernet interfaces. 10.1.2.4 and 11.6.7.9 are assigned to Ethernet 1. 12.1.1.2 and 13.9.7.5 are assigned to Ethernet 2. In this example, the Pipeline 220 routes between all displayed networks. The Pipeline 220 enables the

host assigned 12.1.1.1 to communicate with the host assigned 13.9.7.4 and the host assigned 10.1.2.3.

The host assigned 12.1.1.1 and the host assigned 13.9.7.4 share a physical cable segment, but cannot communicate unless the Pipeline 220 routes between the 12.0.0.0 network and the 13.0.0.0 network.

Enabling RIP on the Ethernet interface

You can configure each IP interface to send RIP updates (informing other local routers of its routes), receive RIP updates (learning about networks that can be reached via other routers on the Ethernet), or both.

Note: Ascend recommends that you run RIP version 2 (RIP-v2) if possible. You should 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 does not propagate subnet mask information, and the default-class network mask is assumed, while RIP-v2 handles subnet masks explicitly. Running the two versions on the same network can result in RIP-v1 class subnet mask assumptions overriding accurate subnet information obtained via RIP-v2.

Ignoring the default route

You can configure the Pipeline 220 to ignore default routes advertised by routing protocols. This configuration is recommended, because you typically do not want the default route changed by a RIP update. The default route specifies a static route to another IP router, which is often a local router such as an Ascend GRF400 or other kind of LAN router. When the Pipeline 220 is configured to ignore the default route, RIP updates do not modify the default route in the Pipeline 220 routing table.

Proxy ARP and inverse ARP

The Pipeline 220 can be configured to respond to ARP requests for remote devices that have been assigned an address dynamically. It responds to the ARP request with its own MAC address while bringing up the connection to the remote device. This feature is referred to as Proxy ARP.

The Pipeline 220 also supports Inverse Address Resolution Protocol (Inverse ARP). Inverse ARP allows the Pipeline 220 to resolve the protocol address of another device when the hardware address is known. The Pipeline 220 does not issue any Inverse ARP requests, but it does respond to Inverse ARP requests that have the protocol type of IP (8000 hexadecimal), or in which the hardware address type is the two-byte Q.922 address (Frame Relay). All other types are discarded. The Inverse ARP response packet sent by the Pipeline 220 includes the following information:

- ARP source-protocol address is the Pipeline 220 unit's IP address on Ethernet.
- ARP source-hardware address is the Q.922 address of the local DLCI.

For the details of Inverse ARP, see RFCs 1293 and 1490.

Telnet password

The Telnet password is required from all users attempting to access the Pipeline 220 unit via Telnet. Users are allowed three tries to enter the correct password, after which the connection attempt fails.

BOOTP relay

By default, a Pipeline 220 does not relay Bootstrap Protocol (BOOTP) requests to other networks. It relays them if BOOTP is enabled. However, the SLIP BOOTP check box must be cleared in System > Terminal Server > General. SLIP BOOTP makes it possible for a computer connecting to the Pipeline 220 over a SLIP connection to use the Bootstrap Protocol. A Pipeline 220 can support BOOTP on only one connection. If both SLIP BOOTP and BOOTP relay are enabled, you receive an error message.

You can specify the IP address of one or two BOOTP servers. You are not required to specify a second BOOTP server.

Note: If you specify two BOOTP servers, the Pipeline 220 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.

Local domain name

The Pipeline 220 uses the Protocols button > IP tab > DNS button > Domain Name parameter for DNS lookups. When the Pipeline 220 receives a hostname to look up, it tries various combinations including appending the configured domain name. The secondary domain name (Protocols button > IP tab > DNS button > 2nd Domain Name parameter) can specify another domain name that the Pipeline 220 can search through DNS. The Pipeline 220 searches the secondary domain only after the domain specified in the Domain Name parameter.

DNS or WINS name servers

When the Pipeline 220 is informed about DNS (or WINS), Telnet and Rlogin users can specify hostnames instead of IP addresses. If you configure a primary and secondary name server, the secondary server is accessed only if the primary server is inaccessible.

DNS lists

DNS can return multiple addresses for a hostname in response to a DNS query, but it does not include information about availability of those hosts. Users typically attempt to access the first address in the list. If that host is unavailable, the user must try the next host, and so forth. However, if the access attempt occurs automatically as part of immediate services, the physical connection is torn down when the initial connection fails. To avoid tearing down physical links when a host is unavailable, you can set the Protocols button > IP tab > DNS button > Enable List Attempt parameter to allow multiple attempts before terminating the WAN session. The Size parameter specifies the maximum number of hosts listed (up to 35).

Client DNS

Client DNS configurations define DNS-server addresses presented to WAN connections during IPCP negotiation. The configurations provide a way to protect your local DNS

information from WAN users. Client DNS has two levels: a global configuration that applies to all PPP connections (defined in Protocols button > IP tab > DNS button), and a connection-specific configuration that applies only to the WAN connection defined in Connection button > IP tab > DNS button. The global client addresses are used only if none are specified in the connection.

SNTP service

The Pipeline 220 can use Simple Network Time Protocol (SNTP), defined in RFC 1305, to set and maintain its system time by communicating with an SNTP server. SNTP must be enabled for the Pipeline 220 to communicate by means of that protocol. In addition, you must specify your time zone as an offset from the Universal Time Configuration (UTC). UTC is in the same time zone as Greenwich Mean Time (GMT), and the offset is specified in hours, using a 24-hour clock. Because some time zones, such as Newfoundland, cannot use an even-hour boundary, the offset includes four digits and is stated in half-hour increments. For example, in Newfoundland the time is 1.5 hours ahead of UTC, and is represented as follows:

UTC +0130

For San Francisco, which is 8 hours ahead of UTC, you would enter:

UTC +0800

For Frankfurt, which is 1 hour behind UTC:

UTC -0100

Specifying SNTP server addresses

The System button > Date & Time tab > Host parameter lets you specify up to three server addresses. The Pipeline 220 will attempt to communicate with the first address. It will attempt the second only if the first is inaccessible, and the third only if the second is inaccessible.

UDP checksums

If data integrity is of the highest concern for your network, and having redundant checks is important, you can set the Protocols button > IP tab > Generate UDP Checksums parameter to generate a checksum whenever a UDP packet is transmitted. UDP packets are transmitted for queries and responses related to ATMP, SYSLOG, DNS, ECHOSERV, RIP, SNTP, and TFTP.

Selecting UDP checksums could cause a slight decrease in performance, but in most environments the decrease is not noticeable.

Poisoning dialout routes in a redundant configuration

If you have another Ascend unit backing up the Pipeline 220 in a redundant configuration on the same network, you can use the Protocols button > All Protocols tab > Advertise Dialout Routes parameter to instruct the Pipeline 220 to stop advertising IP routes if its trunks are in the alarm condition. Otherwise, it continues to advertise its dialout routes, which prevents the redundant unit from taking over the routing responsibility.

Examples of IP network configurations

This section shows an example of a simple system-IP configuration for the Ethernet interface of the Pipeline 220, and a more complete example with system, route, and connection configurations that work together.

Configuring the Pipeline 220 IP interface on a subnet

On a large corporate backbone, many sites 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 it supports an Ascend GRF router at 10.0.0.17, as shown in Figure 6-5.




Figure 6-5. Creating a subnet for the Pipeline 220

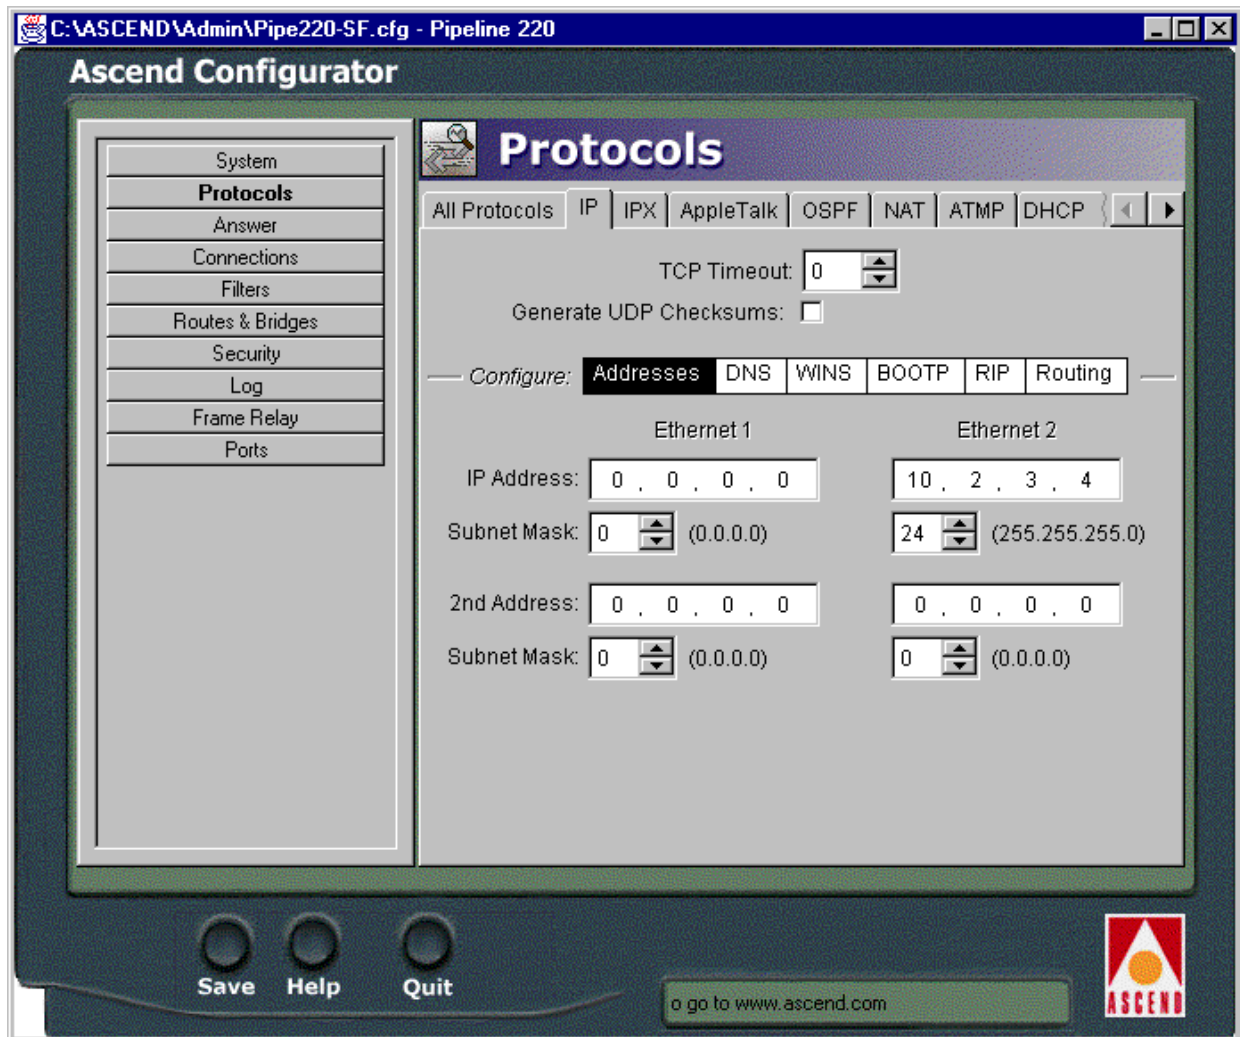
Joining a subnet

You can place the Pipeline 220 on a subnet of that network by entering a subnet mask in its IP address specification. For example:

- 1 On the left side of the Configurator, click the Protocols button.
The right side of the Configurator now displays protocol tabs.
- 2 Click the IP tab.
IP-option buttons appear in the right section of the Configurator.

- 3 Click the Addresses button.

IP-address parameters appear in the lower-right section of the Configurator:



- 4 Click in the IP Address field, and enter a value of 10.2.3.1 for the Pipeline 220.
- 5 Click in the Subnet mask field, and enter a value of 24.
- 6 Click the RIP tab.
- 7 Click in the RIP field, and select Receive-v2.

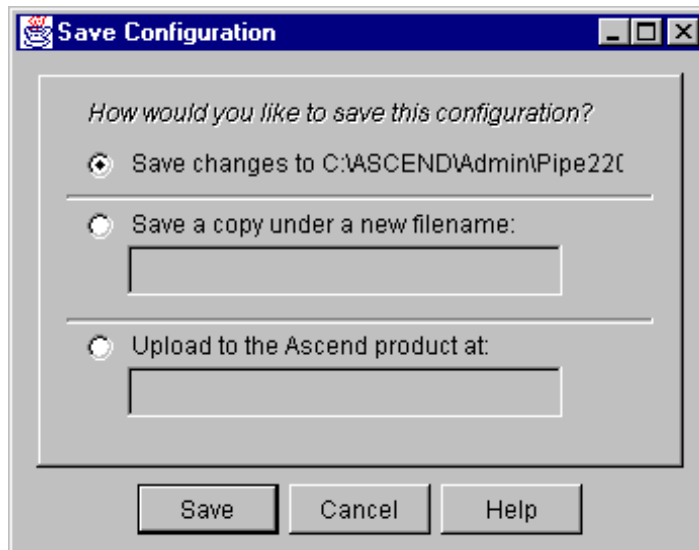
This is an optional step that configures the Pipeline 220 to receive RIP updates from the local router.

- 8 Enable the new settings, as described in “Saving the settings” on page 6-13.

Saving the settings

The Pipeline 220 settings you configure with the Ascend Configurator do not take effect until you upload them to the Pipeline 220. You use the Save command to upload the changes. Alternatively, you can save the settings to a text file for subsequent use, or to another device. To save the settings:

- 1 In the lower-left corner of the Configurator, click Save.
A dialog box appears, prompting you for a save method:



- 2 Select one of the following:
 - If you are uploading the configuration to the same device, select "Upload changes to <Pipeline 220 name>."
 - If you are saving the configuration to a text file on your computer, select "Save a copy under a new filename" and enter the filename to which the configuration should be saved.
 - If you are uploading the configuration to a different device, select "Upload to the Ascend product at:" and enter the IP address of the device.
- 3 Click Save.

If you selected "Upload changes to <Pipeline 220 name>" in step 2, the Ascend

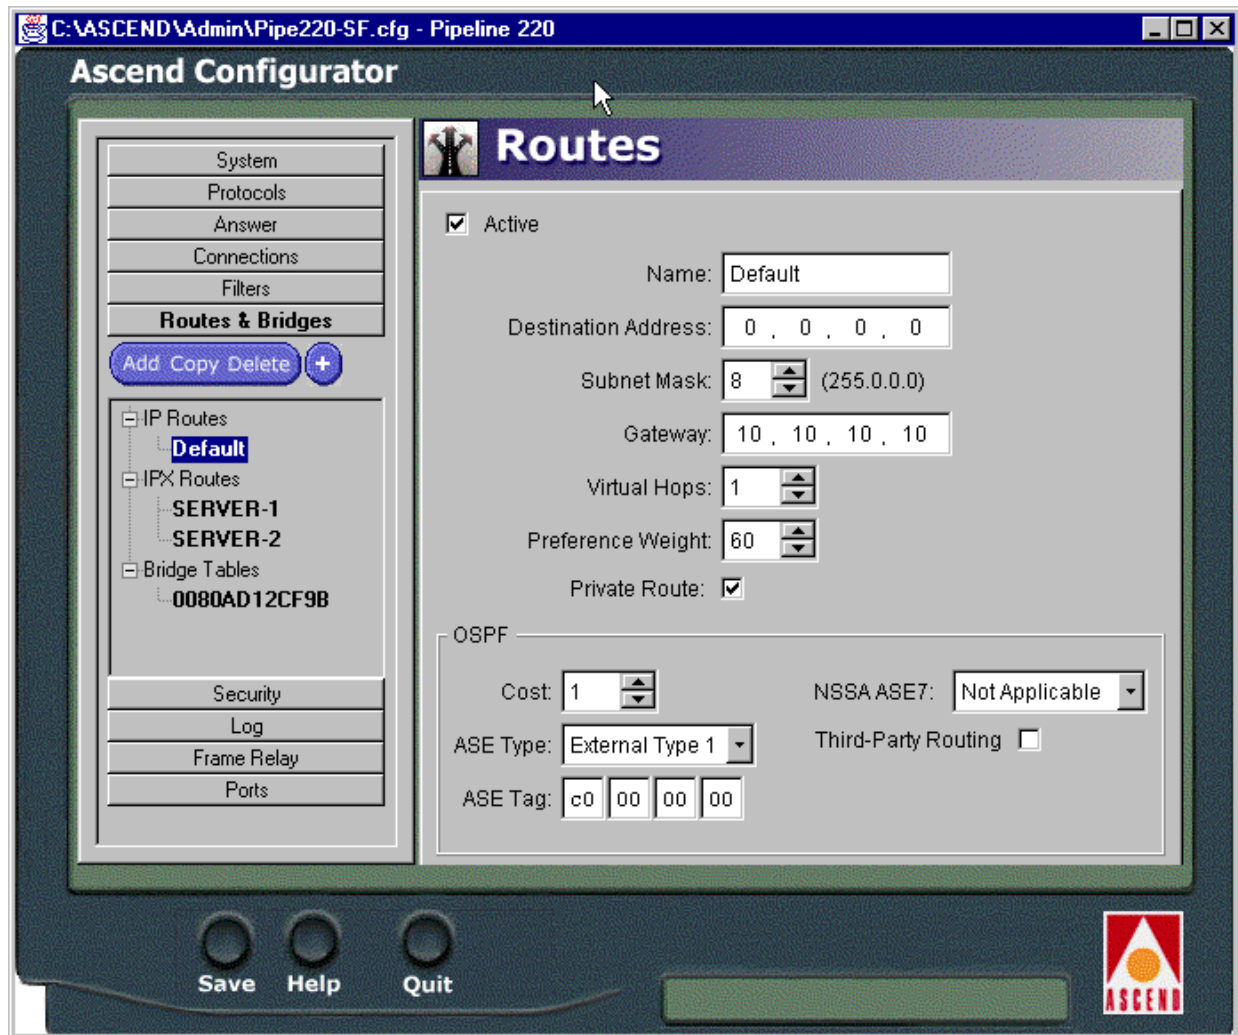
Configurator uploads the configuration file to the Pipeline 220 via TFTP. If you selected one of the other options, the Configurator sends the file to the location you specified. If you specified a remote device, transmission is via TFTP.

Making the backbone router the default route

With the subnet address shown in Figure 6-5, the Pipeline 220 requires a static route to the backbone router on the main network. Otherwise, it can only communicate with devices on the subnets to which it is directly connected. To create the static route and make the backbone router the default route:

- 1 Open the Default IP Route profile.

- 2 Display the Default route by clicking the Routes menu:



- 3 Select the Active check box.
- 4 Specify the IP address of a backbone router in the Gateway field. For example:
Name=Default
Destination Address=0.0.0.0
Subnet Mask=0
Gateway=10.0.0.17
Virtual Hops=1
Preference Weight=100
- 5 Select the Private Route check box.
Typically, default routes should not be advertised to other routers. They are designed for the internal use of the specific router on which they are configured.
- 6 Enable the new settings, as described in “Saving the settings” on page 6-13.

Configuring DNS

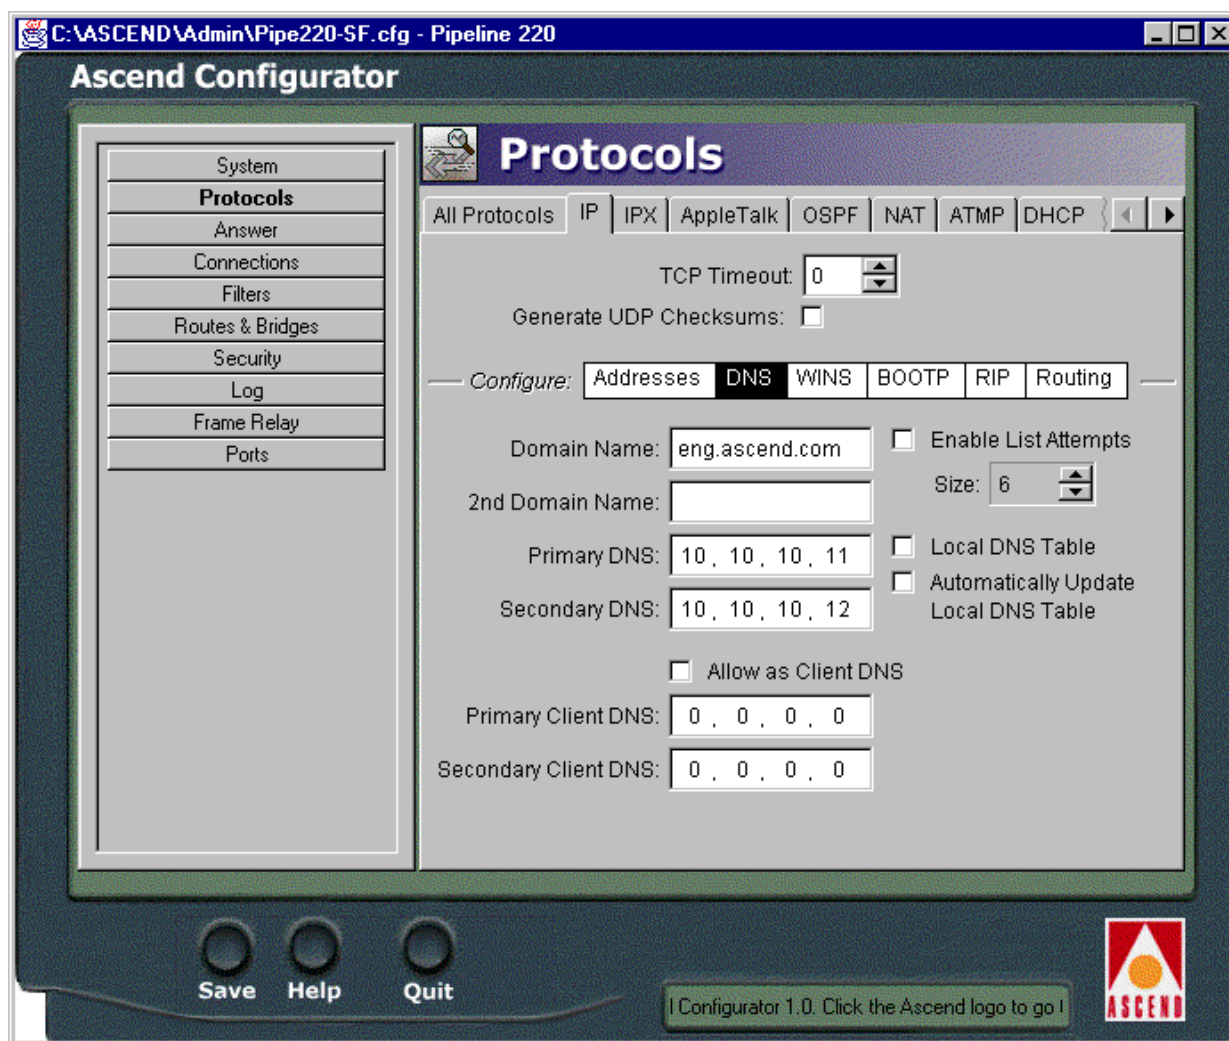
The DNS configuration enables the Pipeline 220 to use local DNS or WINS servers for lookups. In this example of DNS configuration, client DNS is not in use. Note that you can

Configuring IP Routing

Configuring the local IP network setup

protect your DNS servers from callers by defining connection-specific (*client*) DNS servers and specifying that Connection profiles use those client servers. To configure the local DNS service:

- 1 On the left side of the Configurator, click the Protocols button.
The right side of the Configurator now displays protocol tabs.
- 2 Click the IP tab.
IP-option buttons appear in the right section of the Configurator.
- 3 Click the DNS button.
DNS parameters appear in the lower-right section of the Configurator:



- 4 Specify the local Domain Name.
- 5 If appropriate, specify 2nd Domain Name.
- 6 Specify the IP addresses of a primary and secondary DNS server.
- 7 Select the Enable List Attempts and Allow As Client DNS check boxes.

Domain Name=abc.com
2nd Domain Name=
Primary DNS=10.65.212.10


```
Secondary DNS=12.20 7.23.51
Primary Client DNS=0.0.0.0
Secondary Client DNS=0.0.0.0
```

- 8 Enable the new settings, as described in “Saving the settings” on page 6-13.

Configuring IP routing connections

When you enable IP routing and specify addresses in a Connection profile, you define an IP WAN interface. You must configure parameters by clicking both the Answer button and the Connections button. Parameters located under the Answer button enable you to configure parameters for all connections. Parameters located under the Connections button enable you to configure specific values for specific users.

In addition to configuring the Pipeline 220, you should make sure that remote hosts are properly configured.

Understanding the IP routing connection parameters

This section provides some background information about enabling IP routing in the Answer profile and Connection profiles. (For more detailed descriptions of the parameters, see the *Configurator Online Help*.)

Enabling IP routing for WAN connections

You must select the Route IP check box in the Answer button > Routing tab to enable the Pipeline 220 to negotiate a routing connection.

Enabling IP routing for a WAN interface

To enable IP packets to be routed for the connection you are configuring, select the Connections button > IP tab > Enable IP Routing check box. With IP routing enabled, IP packets are always routed, never bridged.

Configuring the remote IP address

The Protocols button > IP tab > Addresses button > IP Address parameter specifies the IP address of the remote device’s Ethernet interface. Before accepting a call from the far end, the

Pipeline 220 matches this address to the source IP address presented by the calling device. It may be one of the following values:

Value	How to specify
IP address of a router	If the remote device is an IP router, specify the address of its Ethernet interface, including its subnet mask modifier. (For background information, see “IP addresses and subnet masks” on page 6-1.) If you omit the subnet mask, the Pipeline 220 inserts a default subnet mask that might make the entire far-end network accessible.
IP address of a remote host	If the remote device is a dial-in host running PPP software, specify its address, including a subnet mask modifier of 32.
The null address (0.0.0.0)	If the remote device is a dial-in host that will accept dynamic address assignment, leave the remote-address parameter blank.

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

WAN alias

A WAN alias is another IP address for the remote device, used for numbered-interface routing. The WAN Alias is listed in the routing table as a gateway (next hop) to the IP Address. The caller must use a numbered interface, and its interface address must agree with the WAN Alias setting.

Specifying a local IP interface address

The Connections button > IP tab > Addresses button > Interface Address is the Pipeline 220's IP address on its WAN interface. In some environments, routing might not work correctly if the Pipeline 220 uses the default, which is its Ethernet interface IP Address.

Assigning metrics and preferences

To favor specific links, you can assign higher metrics to less desirable connections to the same location.

Each connection represents a static route, which has a default preference of 100. (For other preferences, see “Route preferences and metrics” on page 6-5.) For each connection, you can fine-tune the route preference or assign a completely different preference.

Private routes

The Connections button > IP tab > Private Address parameter specifies whether the Pipeline 220 will disclose the existence of the route when queried by RIP or another routing protocol. Private routes are used internally. They are not advertised.

Configuring RIP policy on the WAN interface

You can configure each WAN connection to send RIP updates (informing other routers on that interface of its routes), receive RIP updates (learning about distant networks from the remote routers), or both.

Ascend recommends that you run RIP version 2 (RIP-v2) if possible. You should 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 does not propagate subnet mask information, and the default class network mask is assumed, while RIP-v2 handles subnet masks explicitly. Running the two versions on the same network can result in RIP-v1 subnet mask assumptions overriding accurate subnet information obtained via RIP-v2.

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 software

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.

Windows or OS/2 software

PCs running Windows or OS/2 need TCP/IP networking software. The software is included with Windows 95, but you might need to purchase and install it separately if your computer has an older version of Windows or OS/2.

Macintosh software

Macintosh computers need MacTCP or Open Transport software for TCP/IP connectivity. MacTCP is included with all Apple system software including and later than Version 7.1. To see if a Macintosh has the software, the user should open the Control Panels folder and look for MacTCP or MacTCP Admin.

Software configuration

For any platform, the TCP/IP software must be configured with the host's IP address and subnet mask. If the host will obtain its IP address dynamically from the Pipeline 220, the TCP/IP software must be configured to allow dynamic allocation. 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 with the Pipeline 220 as its default router.

Examples of IP routing connections

This section provides sample Connection profile configurations for IP routing. These examples all assume that the Ethernet profile has been configured correctly, as described in “Configuring the local IP network setup” on page 6-8.

Configuring a router-to-router connection

In the following example, the Pipeline 220 is connected to a corporate IP network and is to be connected to another company that has its own IP configuration. Figure 6-6 shows the network diagram.




Figure 6-6. A router-to-router IP connection

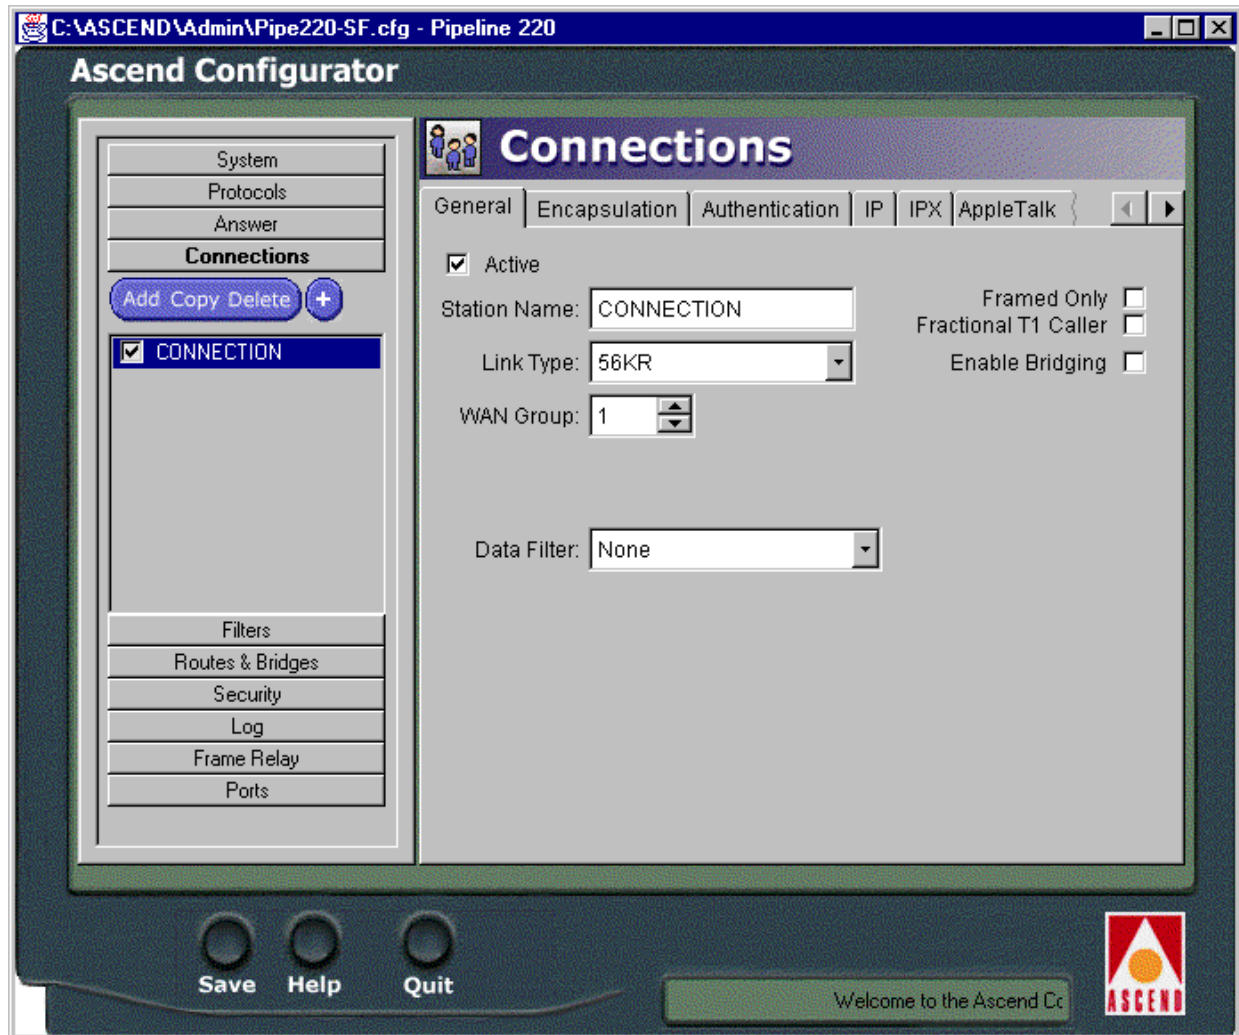
This example assumes that the Answer profile in each of the two devices enables IP routing.

Configuring a Connection profile to the remote device

To configure the site A Pipeline 220 for a connection to remote device B:

- 1 On the left side of the Configurator, click the Connections button.
The Add/Copy/Delete window appears in the lower-left side of the Configurator.
- 2 In the Add/Copy/Delete window, click Add.
The Configurator displays a dialog box prompting you for a name of the new connection.

- 3 Enter a name for the connection, then click OK. For example:
Station Name=CONNECTION
The Connection profile appears on the right side of the Configurator, displaying the name and all other values at their factory default settings:



- 4 On the right side of the Configurator, click the General tab.
- 5 Select the Active check box and set the Station Name parameter to the name of the remote device. For example:
Station Name=PipelineB
- 6 Click the Encapsulation tab and set encapsulation options.
- 7 Click the IP tab to configure IP routing options.
- 8 Select the Enable IP Routing check box.
- 9 Click the Addresses IP-Options button.
- 10 Set the IP Address parameter to the address the remote device's Ethernet interface:
IP Address=10.9.8.10
Subnet Mask=22

Configuring authentication

To configure authentication:

- 1 On the left side of the Configurator, click the Connections button.
- 2 Click the Authentication tab.
- 3 Select the method of authentication. For example:
`Authentication = CHAP`
- 4 Set Dial-out Password to the value sent to the remote device to authenticate the Pipeline 220.
- 5 Set Dial-in Password to the value the remote device sends to the Pipeline 220 to authenticate itself.
- 6 Enable the new settings, as described in “Saving the settings” on page 6-13.

Configuring a router-to-router connection on a subnet

In this sample network, the Pipeline 220 connects telecommuters with their own Ethernet networks to the corporate backbone. The Pipeline 220 is on a subnet, and assigns subnet addresses to the telecommuters' networks.

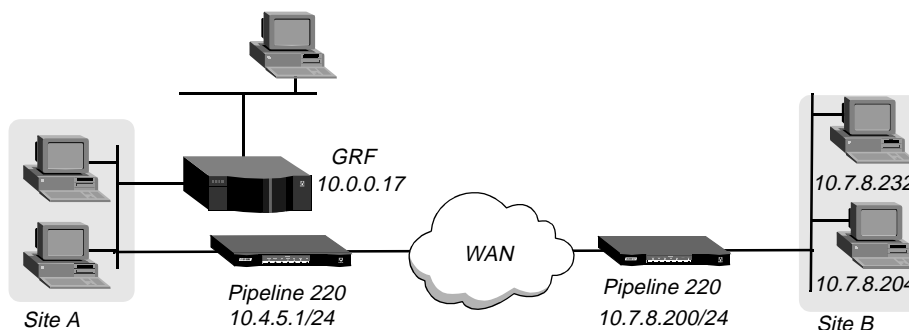


Figure 6-7. A connection between local and remote subnets

This example assumes that both of the Pipeline 220 unit's Answer profiles enable IP routing. Because the Pipeline 220 specifies a subnet mask as part of its own IP address, it must use other routers to reach IP addresses outside that subnet. To forward packets to other parts of the corporate network, the Pipeline 220 either must have a default route configuration to a router in its own subnet or must enable RIP on Ethernet.

Configuring a Connection profile to the remote device

To configure the Pipeline 220 at site A with an IP routing connection to site B:

- 1 On the left side of the Configurator, click the Connections button, and add a Connection profile for the site B device.
- 2 Click the General tab. Select the Active check box and set the Station Name parameter to the name of the remote device. For example:
`Station Name=PipelineB`

- 3 Click the Encapsulation tab and set encapsulation options.
- 4 Click the IP tab to configure IP routing options.
- 5 Select the Enable IP Routing check box.
- 6 Click the Addresses IP-Options button.
- 7 Set the IP Address parameter to the address of the Ethernet interface of the remote device.
IP Address=10.7.8.200
Subnet Mask=24

Configuring authentication

To configure authentication:

- 1 On the left side of the Configurator, click the Connections button.
- 2 Click the Authentication tab.
- 3 Select the method of authentication. For example:
Authentication=CHAP
- 4 Set Dial-out Password to the value sent to the remote device to authenticate the Pipeline 220.
- 5 Set Dial-in Password to the value the remote device sends to the Pipeline 220 to authenticate itself.

Specifying the Pipeline 220's default route

To specify the local backbone router as the Pipeline 220's default route:

- 1 On the left side of the Configurator, click the Routes button.
- 2 Select the Active check box.
- 3 Specify the backbone router's address as the gateway address.
Gateway=10.0.0.17
- 4 Select the Private Route check box.
Typically, default routes should not be advertised to other routers. They are designed for the internal use of the specific router on which they are configured.
- 5 Save the new settings, as described in "Saving the settings" on page 6-13.

Next, configure the remote router for connection to the Pipeline 220. Some important steps you must remember are:

- Make sure the names and passwords are correct for bi-directional authentication. The WAN link will not successfully come up if authentication fails for either the Pipeline 220 or the remote device.
- Double-check the IP addresses to make sure they are correct.
- Be sure to configure the default route of the remote device to the IP address of the Pipeline 220's Ethernet interface.

Configuring a numbered interface

If you are not familiar with numbered interfaces, see “Numbered interfaces” on page 6-6. In the following example, the Pipeline 220 is a system-based router but supports a numbered interface for one of its connections. Figure 6-8 shows an environment, that includes numbered interfaces:




Figure 6-8. Example numbered interface

The numbered interface addresses for the Pipeline 220 are:

- Interface Address=10.5.6.7/24
- WAN Alias=10.5.6.8/24

To configure the numbered interface, first verify the IP address of the Pipeline 220, then configure the WAN connection.

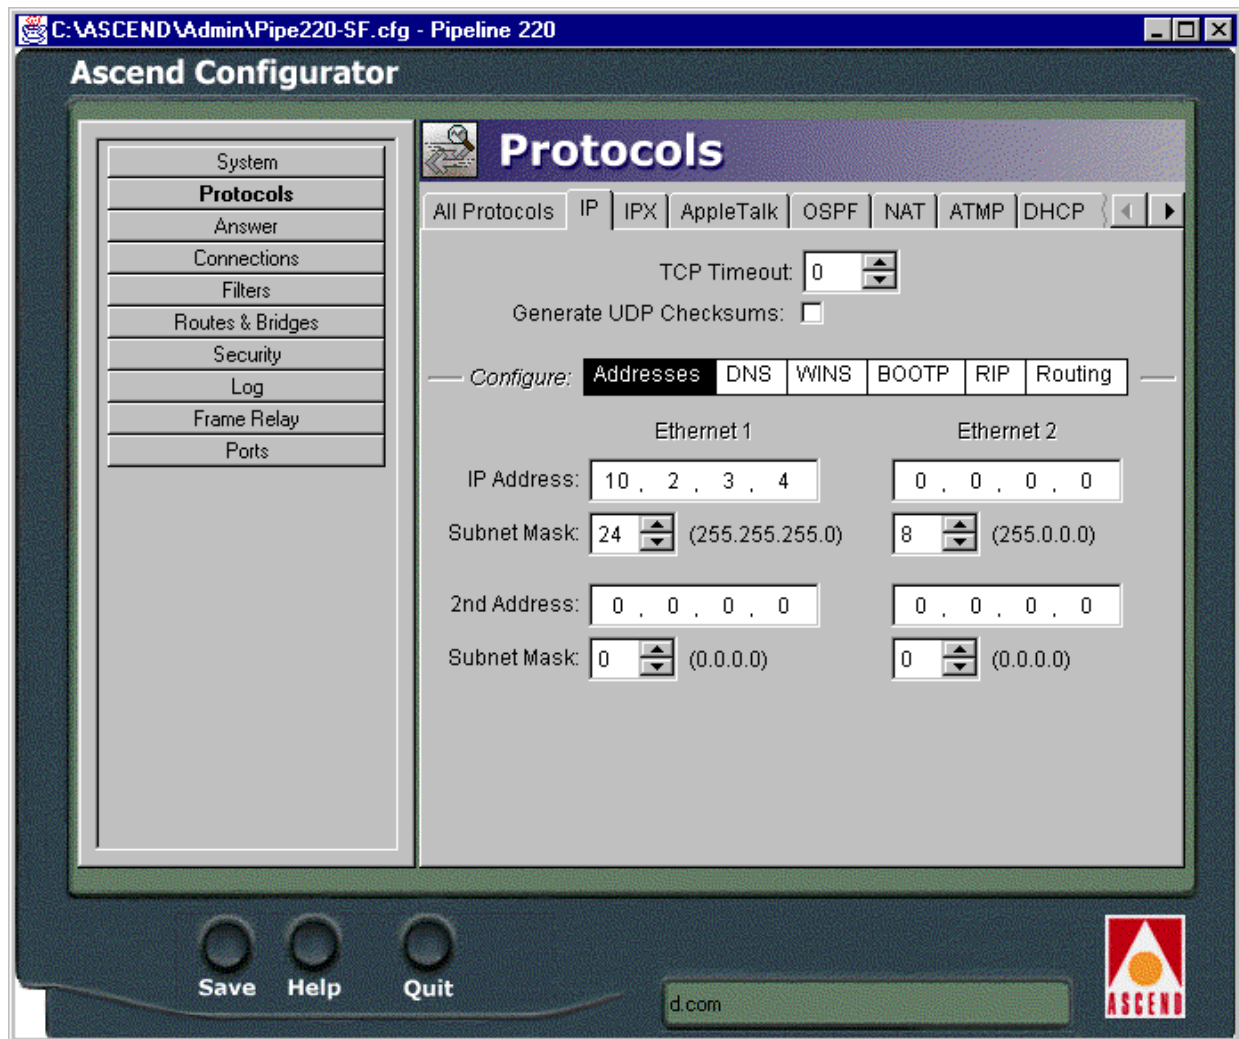
Verifying the IP address of the Pipeline 220

To verify the IP address of the Pipeline 220:

- 1 On the left side of the Configurator, click the Protocols button.
- 2 Click the IP tab.

- 3 Click the Addresses button.

The IP-address parameters appear in the lower-right section of the Configurator:



- 4 Verify that the IP Address and Subnet Mask parameters are set correctly for the Ethernet interface of the Pipeline 220:

IP Adrs=10.2.3.4
Subnet Mask=24

Configuring the WAN connection

To configure the WAN connection and its addresses:

- 1 On the left side of the Configurator, click the Connections button.
The Add/Copy/Delete window appears in the lower-left side of the Configurator.
- 2 In the Add/Copy/Delete window, click Add.
The Configurator displays a dialog box prompting you for a name of the new connection.
- 3 Enter name for the connection, then click OK. For example:
Station Name=PipelineB

Configuring IP Routing

Configuring IP routing connections

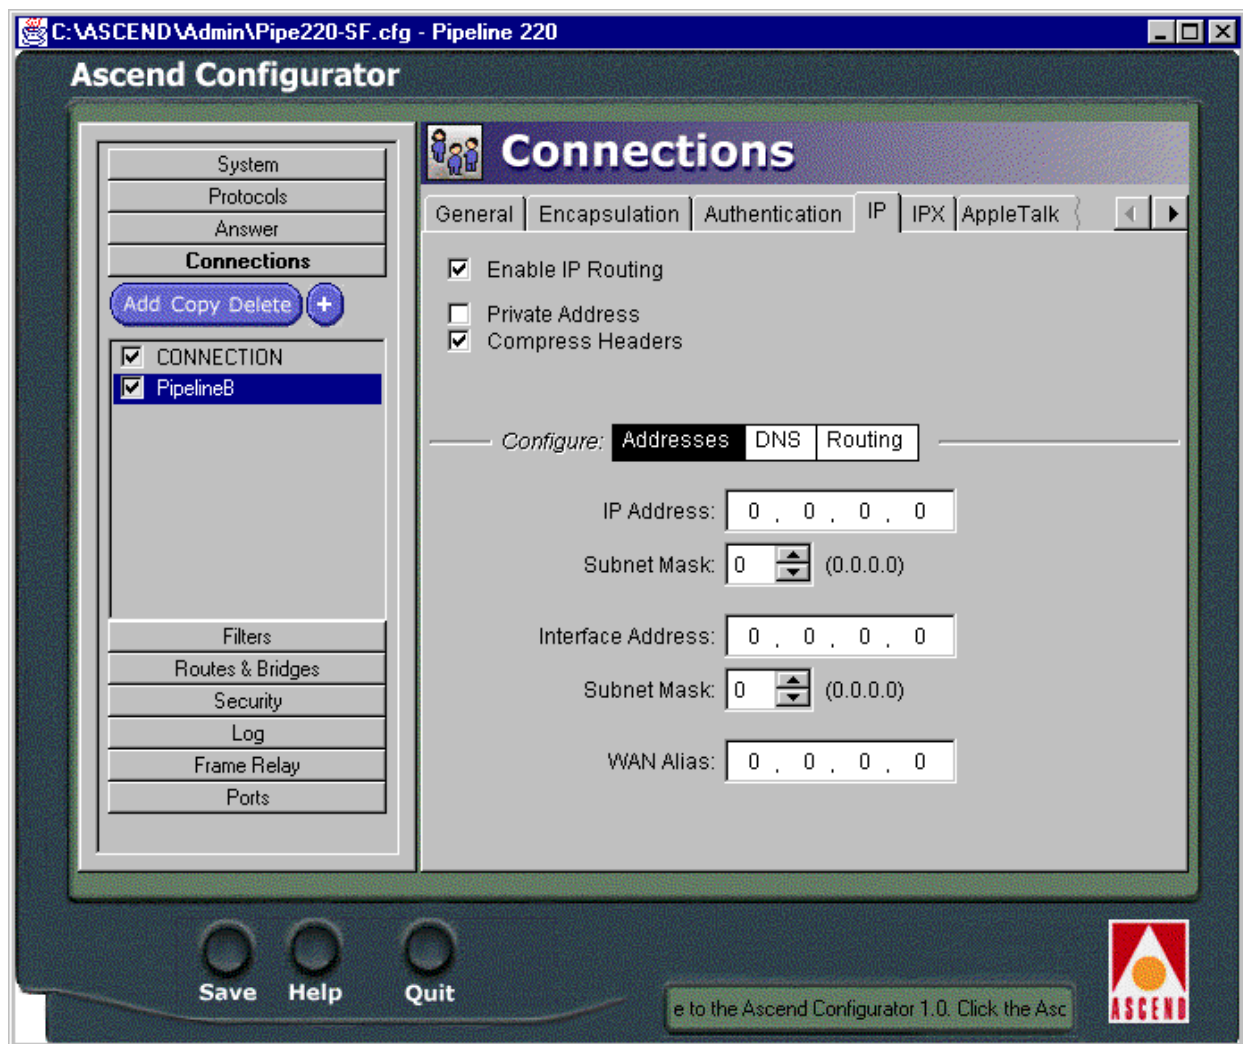
The Connection profile appears on the right side of the Configurator, displaying the name and all other values at their factory default settings.

Make sure that you set the Station Name parameter to the name of the remote device's name. Station Name is used for link authentication. For example:

Station Name=PipelineB

- 4 Click the Encapsulation tab and set encapsulation options.
- 5 Click the IP tab.
- 6 Select the Enable IP Routing check box.
- 7 In the right section of the Configurator, click the Addresses button.

The IP-address parameters appear in the lower-right section of the Configurator:



- 8 Set the IP Address parameter to the address of the Ethernet interface of the remote device.
IP Address=10.3.4.5
Subnet Mask=24
- 9 Set the IP addresses for the WAN link:
WAN Alias=10.7.8.9
Subnet Mask=24

Interface Address=10.5.6.7
Subnet Mask=24

Configuring authentication

To configure authentication:

- 1 On the left side of the Configurator, click the Connections button.
- 2 Click the Authentication tab.
- 3 Select the method of authentication. For example:
`Authentication = CHAP`
- 4 Set Dial-out Password to the value sent to the remote device to authenticate the Pipeline 220.
- 5 Set Dial-in Password to the value the remote device sends to the Pipeline 220 to authenticate itself.
- 6 Enable the new settings, as described in “Saving the settings” on page 6-13.

Configuring IP routes and preferences

The IP routing table contains routes that are configured (static routes) and routes that are learned dynamically from routing protocols such as RIP or OSPF. The section discusses static routes.

Understanding the static route parameters

This section provides some background information on static routes. For complete information about each parameter, see the *Configurator Online Help*.

Route names

IP routes are indexed by name. You can assign any name of less than 31 characters.

Activating a route

A route must be active to affect packet routing. An inactive route is ignored. You activate a route by selecting its Active check box.

Route's destination address

The Routes button > Destination Address of a route is the target network (the destination address in a packet). Packets destined for that host use this static route to bring up the right connection. The zero address (0.0.0.0) represents the default route (the destination to which packets are forwarded when there is no route to the packet's destination).

Route's gateway address

The Routes button > Gateway parameter specifies the IP address of the router or interface by which to reach the target network.

Virtual hops, costs, and preferences

The Routes button > Virtual Hops parameter is a metric or hop count for the current route (a number between 1 to 15). When RIP was originally developed, the hop count was a number that showed how many routers had to be crossed to reach the destination. For example, a destination with a hop count of 10 meant that getting a packet there required going through 10 routers. A route with a shorter hop count to a destination is more desirable than one with a larger hop count, since it most likely is a shorter, faster route.

The hop count can also be manually configured to give a route a *virtual* hop count. In this way, you manually configure which routes are more desirable than others in your environment. With a choice of two identical routes but different hop count, the Pipeline 220 uses the route with the lower hop count.

The Routes button > Cost parameter specifies the cost of an OSPF link. The cost is a configurable metric that takes into account the speed of the link and other issues. The lower the cost, the more likely the interface will be used to forward data traffic. (For details, see Chapter 8, “Configuring OSPF Routing.”)

The Routes button > Preference Weight parameter specifies a route preference. Zero is the default for connected routes (such as for the Ethernet interface). When choosing which route to use, the router first compares the preference values, preferring the lowest number. If the preference values are equal, the router compares the Virtual Hops values, using the route with the lowest virtual hop count. A Preference Weight of 255 means “Don't use this route.” (For more information, see “Route preferences and metrics” on page 6-5.)

Tagging routes learned from RIP

The Routes button > ASE Tag value is attached to all routes learned from RIP in OSPF updates. The tag is a hexadecimal number that can be used by border routers to filter the record.

Type-1 or type-2 metrics for routes learned from RIP

The Routes button > ASE Type parameter can be set to 1 or 2. Type-1 is a metric expressed in the same units as the link-state metric (that is, the same units as interface cost). Type-2 is considered larger than any link-state path. It assumes that routing between autonomous systems is the major cost of routing a packet, and eliminates the need for conversion of external costs to internal link-state metrics.

Making a route private

Private routes are used internally but are not advertised. To make a route private, select the Routes button > Private Route check box.

Note: Typically, default routes should not be advertised to other routers. They are designed for the internal use of the specific router on which they are configured.

Routes for Connection profile interfaces

When an IP routing connection is brought up, the Pipeline 220 activates the route for that WAN interface. The Destination for the route is the remote device's IP address, and the metric and preference values are specified in the Connection profile. If the profile uses a numbered interface, an additional route is created for that interface.

A connected route for the Ethernet IP interface

The Protocols button > IP tab > Addresses > IP Address parameter specifies the Pipeline 220's IP address on the local Ethernet interface. Remember that the Pipeline 220 has two Ethernet ports, so it has two local Ethernet interfaces. The Pipeline 220 creates a route for these addresses at system startup.

Static route preferences

By default, static routes and RIP routes have the same preference, so they compete equally. ICMP redirects take precedence over both, and OSPF routes take precedence over everything. If a dynamic route's preference is lower than that of the static route, the dynamic route can replace, or *hide*, a static route to the same network. This can be seen in the IP routing table, which will have two routes to the same destination. The static route has an *h* flag, indicating that it is hidden and inactive. The active, dynamically learned route is also in the routing table. However, dynamic routes age, and if no updates are received, they eventually expire. In that case, the hidden static route reappears in the routing table.

RIP and OSPF preferences

Because OSPF typically involves a complex environment, its router configuration is described in a separate chapter. See Chapter 8, "Configuring OSPF Routing."

Examples of static route configurations

This section shows two samples of static route configurations. The Pipeline 220 forwards to the *default* route any packet that is destined for an unknown address. Following a discussion of the default route is a discussion of normal static route configuration.

Configuring the default route

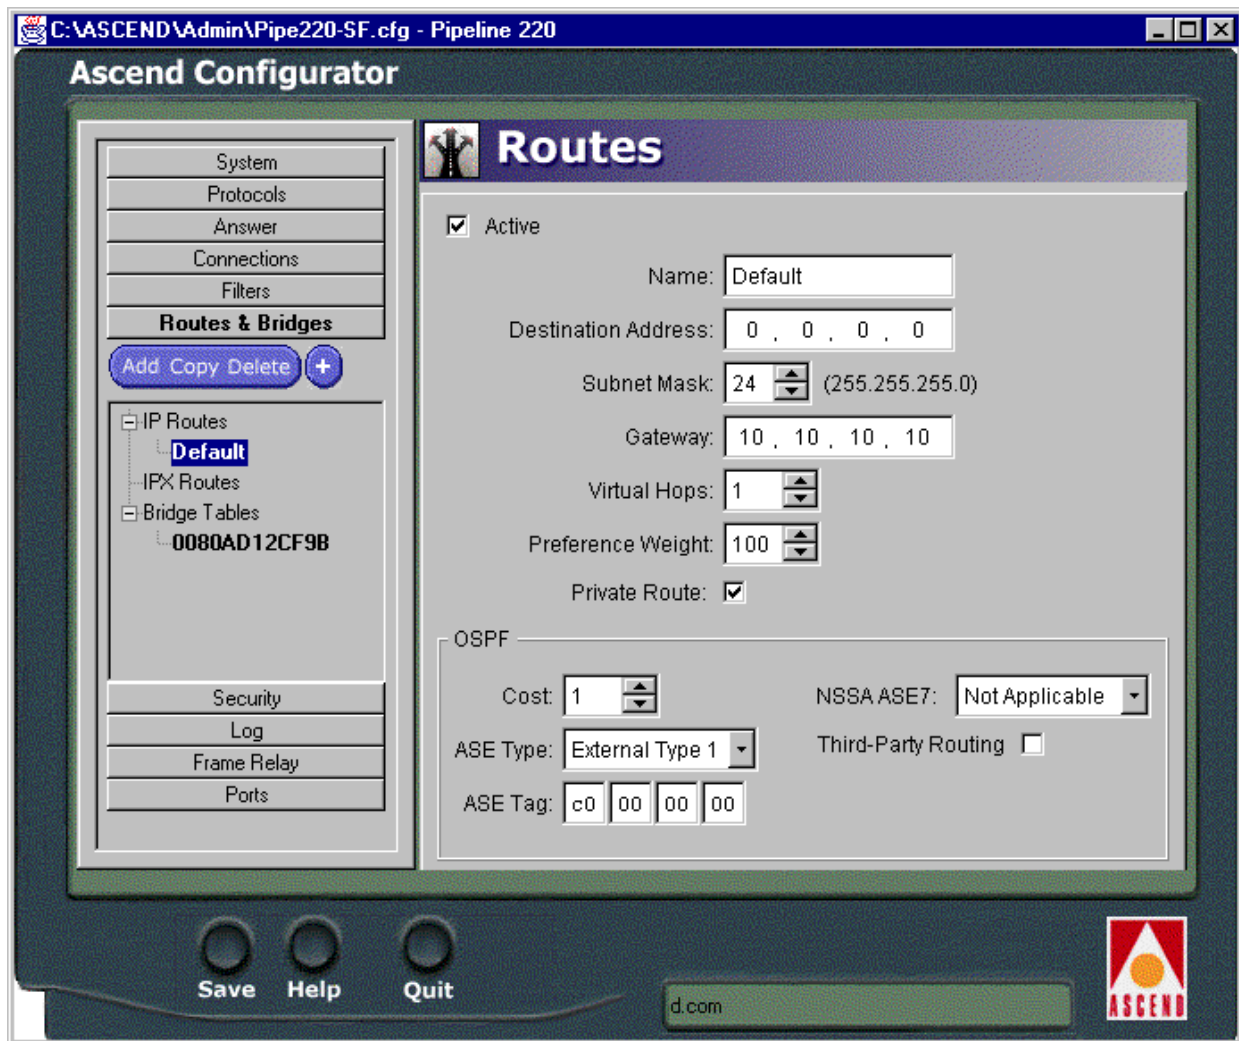
If no routes exist for the destination address of a packet, the Pipeline 220 forwards the packet to the default route. Most sites use the default route to specify a local IP router (such as a Cisco router or a UNIX host running the route daemon) to offload routing tasks to other devices.

Note: If the Pipeline 220 does not have a default route, it drops packets for which it has no route.

The name of the default IP Route profile is always Default, and its destination is always 0.0.0.0. To configure the default route:

- 1** On the left side of the Configurator, click the Routes button.

- 2 From the list of routes on the left side of the Configurator, highlight Default. Specific information about the Default route appears on the right:



The name of the first IP Route profile is always Default, and its destination is always 0.0.0.0. You cannot change these values.

- 3 To activate the route, select the Active check box.
- 4 Specify the router to use for packets with unknown destinations. For example:
Gateway=10.9.8.10
- 5 Specify a metric or hop count for this route and the route's preference. For example:
Virtual Hops=1
Preference Weight=100
- 6 Select the Private Route check box.
Typically, default routes should not be advertised to other routers. They are designed for the internal use of the specific router on which they are configured.
- 7 Enable the settings, as described in "Saving the settings" on page 6-13.

Defining a static route to a remote subnet

If the connection does not enable RIP, the Pipeline 220 does not learn about other networks or subnets that are reachable through the remote device, such as the remote network shown in Figure 6-9.





Figure 6-9. Two-hop connection that requires a static route when RIP is off

To enable the Pipeline 220 to route to site C without using RIP, you must configure a static IP route profile:

- 1 On the left side of the Configurator, click the Routes button.
- 2 From the option bar for the list of routes on the left side of the Configurator, click Add.

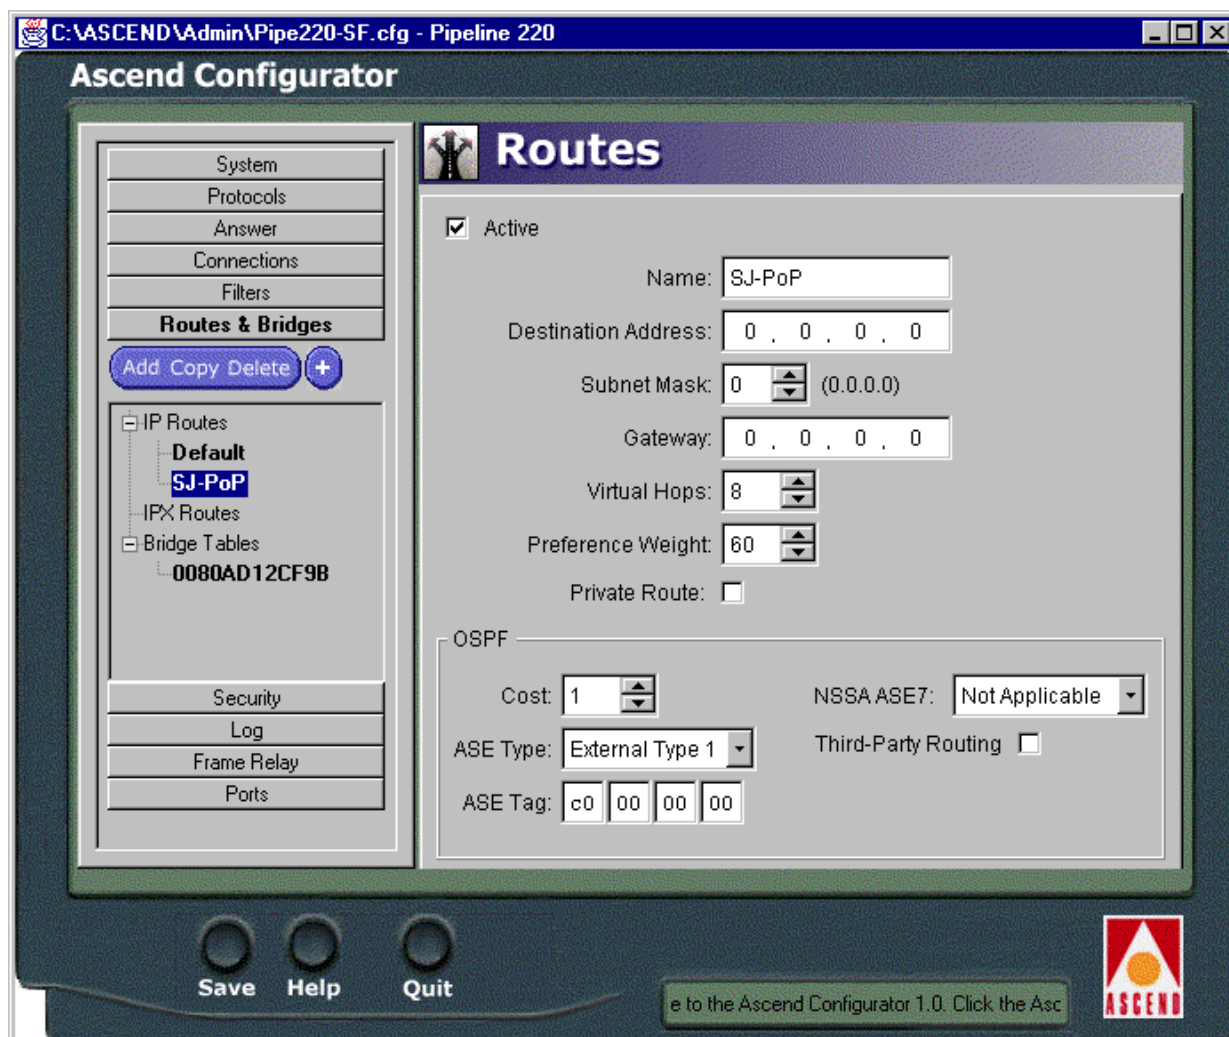
A dialog box appears:



- 3 Enter a name for the route.

The name you choose does not affect routing. It allows you to use descriptive names for routes, in addition to their numerical representations.

The Route profile appears on the right side of the Configurator, displaying the name and all other values at their factory default settings:



- 4 To activate the route, select the Active check box.
- 5 Specify the destination network for this route. For example:
Destination Address=10.4.5.0
Subnet Mask=22
- 6 Specify the router to use for packets destined for the network specified above. For example:
Gateway=10.9.8.10
- 7 Specify a metric or hop count for this route and the route's preference. For example:
Virtual Hops=2
Preference Weight=100
- 8 If you use OSPF, specify values for OSPF-related parameters. For example:

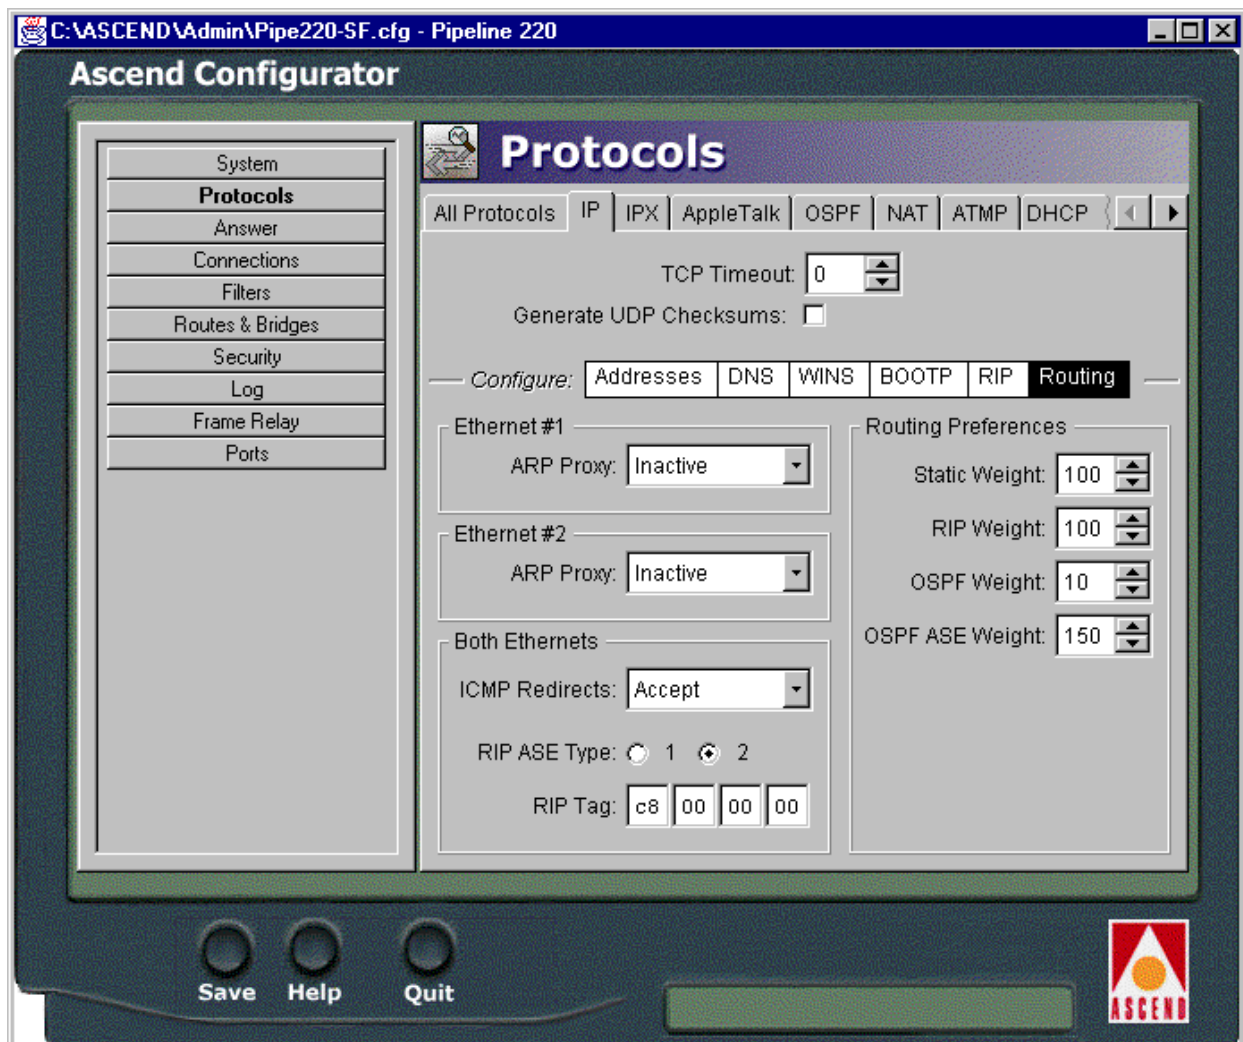

```
Cost=1
ASE Type=External Type1
ASE Tag=c0 00 00 00
```

- 9 Enable the new settings, as described in “Saving the settings” on page 6-13.

Example route preferences configuration

The following example increases the preference value of RIP routes, instructing the router to use a static route first if one exists.

- 1 On the left side of the Configurator, click the Protocols button.
- 2 On the right side of the Configurator, click the IP tab.
- 3 In the lower-right section of the Configurator, click the Routing button.
IP-routing parameters appear on the right side of the Configurator:



- 4 Set the Static Weight parameter to 150.
- 5 Enable the new settings, as described in “Saving the settings” on page 6-13.

Configuring the Pipeline 220 for dynamic route updates

Each active interface can be configured to send or receive RIP or OSPF updates. The Ethernet interface can also be configured to accept or ignore ICMP redirects. All of these routing mechanisms modify the IP routing table dynamically.

Understanding the dynamic routing parameters

This section provides some background information about the dynamic routing options. For complete information about each parameter, see the *Configurator Online Help*. (For information about OSPF updates, see Chapter 8, “Configuring OSPF Routing.”)

RIP (Routing Information Protocol)

You can configure the router to send or receive RIP updates (or both) on the Ethernet interface and on each WAN interface. You can also choose between RIP-v1 and RIP-v2 on any interface. Many sites turn off RIP on WAN connections to keep their routing tables from becoming very large.

Note: The IETF has voted to move RIP-v1 into the *historic* category and its use is no longer recommended. Ascend recommends that you upgrade all routers and hosts to RIP-v2. If you must maintain RIP-v1, you should create a separate subnet and place all RIP-v1 routers and hosts on that subnet.

Ignoring the default route

You can configure the Pipeline 220 to ignore default routes advertised by routing protocols. This configuration is recommended, because you typically do not want the default route changed by a RIP update. The default route specifies a static route to another IP router, which is often a local router such as an Ascend GRF 400 or other kind of LAN router. When the Pipeline 220 is configured to ignore the default route, RIP updates do not modify the default route in the Pipeline 220 routing table.

RIP policy and RIP summary

The RIP Policy and RIP Summary parameters have no effect on RIP-v2.

If the Pipeline 220 is running RIP-v1, the RIP Policy parameter specifies a split horizon or poison reverse policy to handle update packets that include routes that are received on the same interface on which the update is sent. Split-horizon means that the Pipeline 220 does not propagate routes back to the subnet from which they were received. Poison-reverse means that it propagates routes back to the subnet from which they were received, with a metric of 16.

The RIP Summary parameter specifies whether to summarize subnet information when advertising routes. If the Pipeline 220 summarizes RIP routes, it advertises a route to all the subnets in a network of the same class. For example, the route to 200.5.8.13/28 (a class C address subnetted to 28 bits) would be advertised as a route to 200.5.8.0. When the Pipeline 220 does not summarize information, it advertises each route in its routing table *as-is*. For the subnet in the preceding example, the Pipeline 220 would advertise a route only to 200.5.8.13.

Ignoring ICMP redirects

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. They are also one of the least secure methods, because it is possible to counterfeit ICMP Redirects and change the way a device routes packets.

Private routes

If you configure a profile with the Private Route check box selected, the router does not disclose its route in response to queries from routing protocols.

Examples of RIP and ICMP configuration

The following sample configuration instructs the router to ignore ICMP Redirect packets, to receive (but not send) RIP updates on Ethernet, and to send (but not receive) RIP updates on a WAN connection.

Configuring RIP policy

To configure RIP policy on the Pipeline 220:

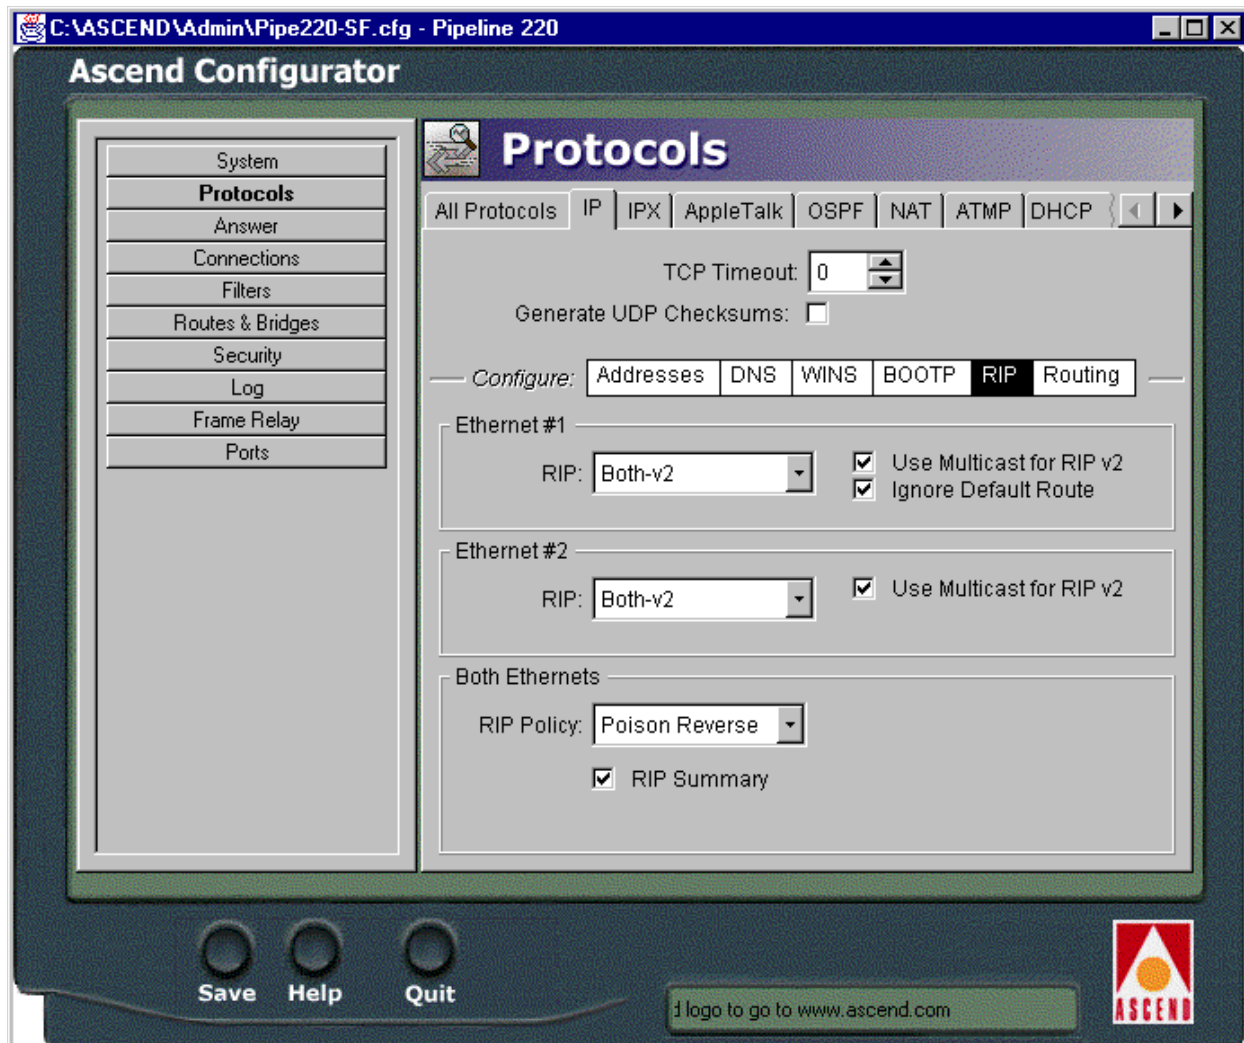
- 1** On the left side of the Configurator, click the Protocols button.
- 2** On the right side of the Configurator, click the IP tab.

Configuring IP Routing

Configuring the Pipeline 220 for dynamic route updates

- 3 On the lower-right side of the Configurator, click the RIP button.

RIP-policy parameters appear in the lower-right section of the Configurator:



- 4 Configure the router to receive (but not send) RIP updates on the Ethernet interface:

RIP=Recv-v2

Receiving RIP updates on Ethernet means that the router will learn about networks that are reachable via other local routers. However, it will not propagate information about all of its remote connections to the local routers.

- 5 Click the Routing button, and set ICMP Redirects to Ignore.

ICMP Redirects=Ignore

Configuring RIP on the WAN link

To configure RIP on the WAN link:

- 1 On the left side of the Configurator, click the Connections button.
- 2 On the right side of the Configurator, click the IP tab.
- 3 After clicking the IP-option buttons on the lower-right section of the Configurator, click the Routing button.

- 4 Configure the router to send (but not receive) RIP updates on this link:

`RIP=Send-v2`

Sending RIP on a WAN connection enables the remote devices to access networks that are reachable via other local routers. However, the Pipeline 220 will not receive information about networks that are reachable through the remote router.

- 5 Enable the new settings, as described in “Saving the settings” on page 6-13.

Syslog services

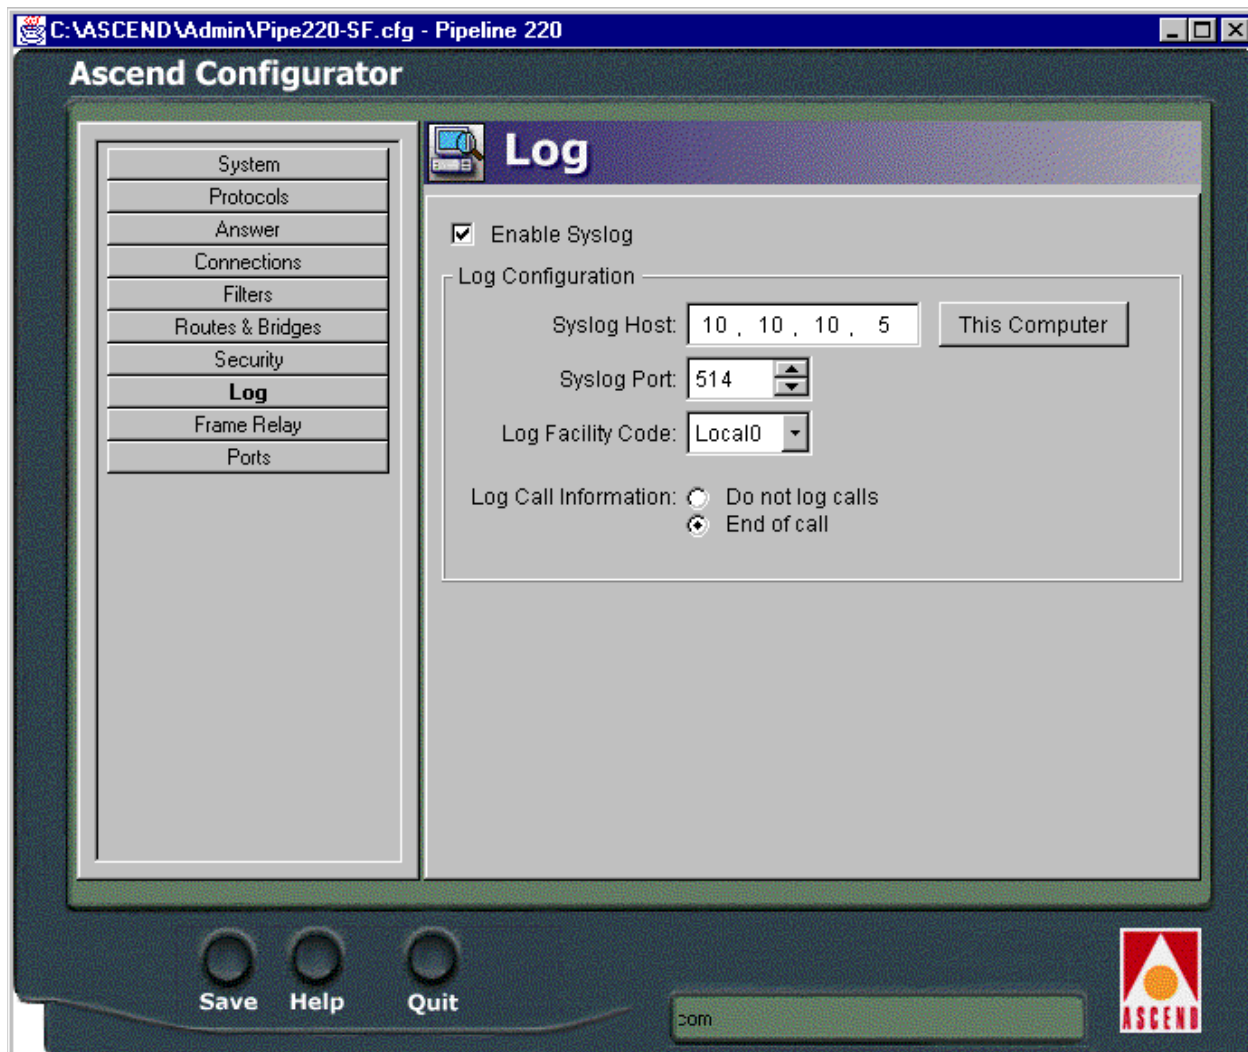
To maintain a permanent log of Pipeline 220 system events and send Call Detail Reporting (CDR) reports to a host that can record and process them, configure the Pipeline 220 to report events to a syslog host on the local IP network. Note that syslog reports are only sent through the Ethernet interface.

Configuring the Pipeline 220 to send Syslog messages

To configure the Pipeline 220 to send messages to a Syslog daemon:

- 1 On the left side of the Configurator, click the Log button.

Syslog parameters appear on the right side of the Configurator:



- 2 Select Enable Syslog.
- 3 Set Syslog Host to the IP address of the host running the Syslog daemon.
If it is the PC running the Configurator, click the “This Computer” button. The Configurator enters the IP address of the PC into the Syslog Host parameter.
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 220, the Pipeline 220 must have a route to that host, either via RIP or a static route.
- 4 Set Syslog Port to the port number at which you want the Syslog host expects to receive messages from this Pipeline 220.
- 5 Set the Log Facility Code.
This parameter is used to flag messages from the Pipeline 220. After you set a Log Facility Code, you must configure the Syslog daemon to write all messages containing that facility code to a particular log file on the Syslog host.

Note: The Log Call Information parameter does not apply to the Pipeline 220. Setting it to either value does not affect operation of the Pipeline 220.

The Pipeline 220 settings you configure with the Ascend Configurator do not take effect until you upload them to the Pipeline 220. You use the Save command to upload the changes. Alternatively, you can save the setting to a text file for subsequent use, or to another device. To save the settings:

- 1** In the lower-left corner of the Configurator, click Save
- 2** A dialog box appears, prompting you for a save method. Select one of the following:
 - If you are uploading the configuration to the same device, select “Upload changes to <Pipeline 220 name>.”
 - If you are saving the configuration to a text file on your computer, select “Save a copy under a new filename” and enter the filename to which the configuration should be saved.
 - If you are uploading the configuration to a different device, select “Upload to the Ascend product at:” and enter the IP address of the device.
- 3** Click Save.

If you selected “Upload changes to <Pipeline 220 name>” in step 2, the Ascend

Configurator uploads the configuration file to the Pipeline 220 via TFTP. If you selected one of the other options, the Configurator sends the file to the location you specified. If you specified a remote device, transmission is via TFTP.

To configure the Syslog daemon, you must modify `/etc/syslog.conf` on the log host. This file specifies which action the daemon performs when it receives messages from a particular log facility number (which represents the Pipeline 220). For example, if you set Log Facility Code to Local5 in the Pipeline 220, and you want to log its messages in `/var/log/Pipeline 220`, add this line to `/etc/syslog.conf`:

```
local5.info<tab>/var/log/Pipeline 220
```

Note: The Syslog daemon must reread `/etc/syslog.conf` after it has been changed.

Syslog messages

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 following standardized 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	<p>The 4 hexadecimal digit Ether Type, or the network protocol name—"arp," "rarp," "ipx," "appletalk."</p> <p>The protocol for IP protocols, is either the IP protocol number (up to 3 decimal digits) or one of the following names:</p> <ul style="list-style-type: none">• 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 the packet is internally inconsistent• unreachable the packet was generated by an “unreach=” rule in the firewall• !pass the packet was blocked by the data firewall• bringup the packet matches the call firewall• !bringup 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>

IP Address Management

This chapter includes the following topics:

BOOTP Relay	7-1
Saving your settings	7-2
DHCP services	7-3
Local DNS host address table	7-6
User-definable TCP connection retry timeout	7-8
Network Address Translation (NAT)	7-9

BOOTP Relay

The Bootstrap Protocol (BOOTP) defines how a computer on a TCP/IP network can obtain its Internet Protocol (IP) address and other startup information 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.

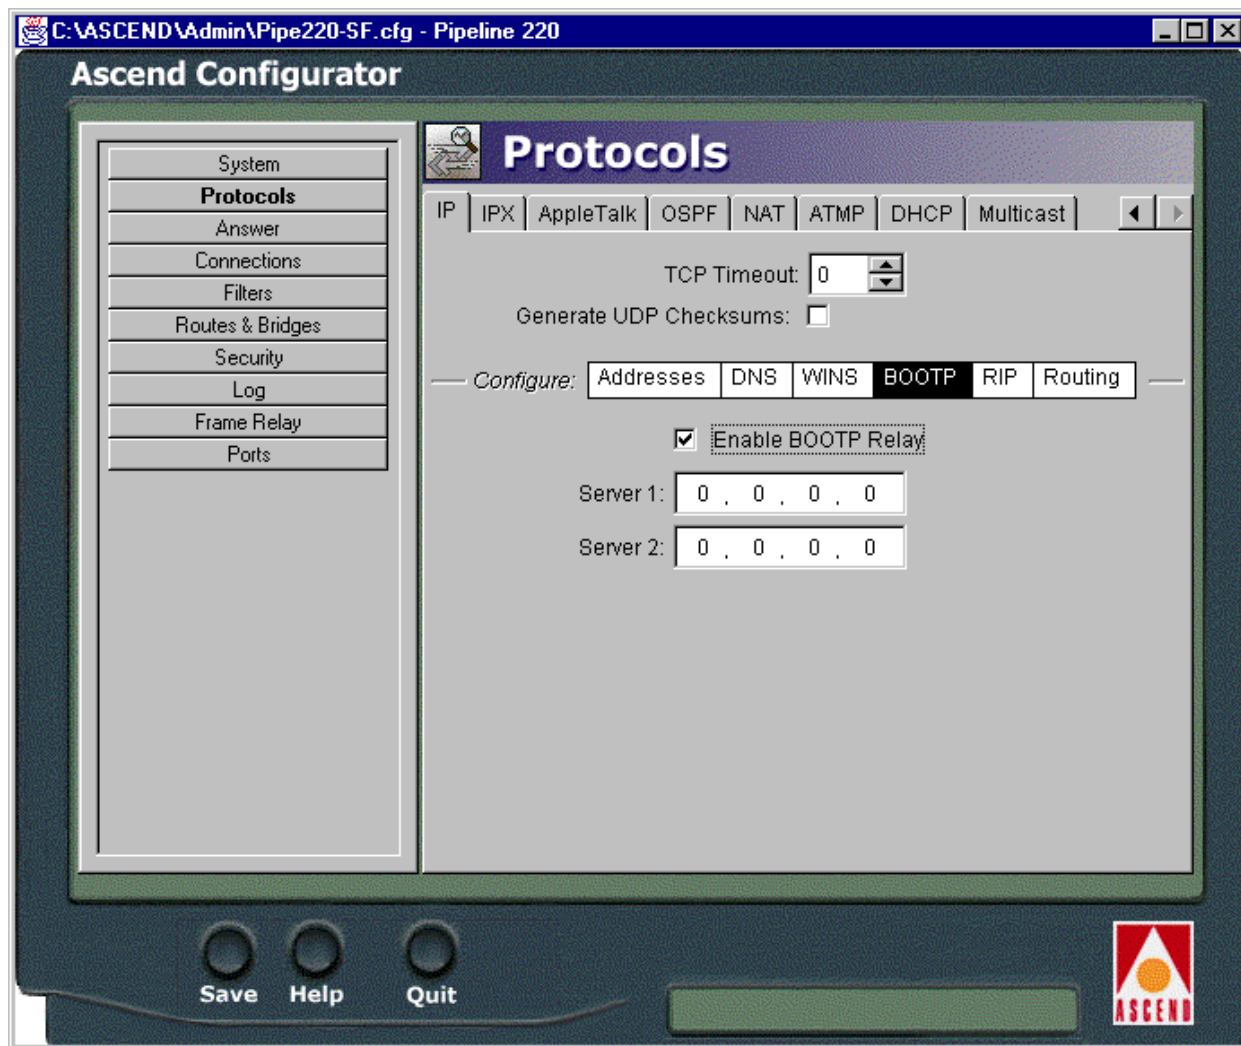
When the BOOTP client and BOOTP server are not on the same local-area network, the BOOTP request must be relayed from one network to another. The Pipeline 220 can perform this task, which is known as BOOTP relay.

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, such as the Pipeline 220, that connects the networks.

By default, a Pipeline 220 does not relay BOOTP requests to other networks. To enable BOOTP relay:

- 1 On the left side of the Configurator, click the Protocols button.
- 2 On the right side of the Configurator, click the IP tab.
- 3 In the lower-right section of the Configurator, click the BOOTP button.

BOOTP parameters appear in the lower-right section of the Configurator:




- 4 Select the “Enable BOOTP Relay” check box.
- 5 Set Server 1 to the IP address of the BOOTP server.
- 6 If there is another BOOTP server available, set Server 2 to its IP address.
If you specify two BOOTP servers, the Pipeline 220 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.
- 7 Enable the new settings, as described in “Saving your settings” on page 7-2.

Saving your settings

The Pipeline 220 settings you configure with the Ascend Configurator do not take effect until you upload them to the Pipeline 220. You use the Save command to upload the changes. Alternatively, you can save the settings to a text file for subsequent use, or to another device. To save the settings:

- 1 In the lower-left corner of the Configurator, click Save.

A dialog box appears, prompting you for a save method:



- 2 Select one of the following:
 - If you are uploading the configuration to the same device, select “Upload changes to *<Pipeline 220 name>*.”
 - If you are saving the configuration to a text file on your computer, select “Save a copy under a new filename” and enter the filename to which the configuration should be saved.
 - If you are uploading the configuration to a different device, select “Upload to the Ascend product at:” and enter the IP address of the device.
- 3 Click Save.

If you selected “Upload changes to *<Pipeline 220 name>*” in step 2, the Ascend Configurator uploads the configuration file to the Pipeline 220 via TFTP. If you selected one of the other options, the Configurator sends the file to the location you specified. If you specified a remote device, transmission is via TFTP.

DHCP services

A Pipeline 220 can perform a number of Dynamic Host Configuration Protocol (DHCP) services, including:

- DHCP Server functions, responding to DHCP requests for up to 46 clients at any given time. DHCP server responses provide an IP address and subnet mask. You can define two address pools, of up to 20 addresses each.
Additionally, you can specify that specific hosts, up to six, receive specific IP addresses. The Pipeline 220 identifies the hosts by their Ethernet (MAC) addresses.
- Management of 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 220's Ethernet interface. The IP address is good for 60 seconds. Once a 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 220 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:

Method	Description
Plug-and-Play	If you select the Protocols button > IP tab > DHCP button > "Enable DHCP PNP" check box, the Pipeline 220 increments its own IP address by one, and returns that address 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 220 to respond to distant networks without having to configure an IP address first.
Renewal	If the host is renewing the address it currently has, the Pipeline 220 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 220 always provides the same address.
Next available	If the host is making a new request and there is no IP address reserved for the host, the Pipeline 220 assigns the next available address from its address pools. The Pipeline 220 assigns the first available address from the first pool IP addresses. If there are no available addresses in the first pool, the Pipeline 220 assigns the first available address from the second pool.

Configuring DHCP services

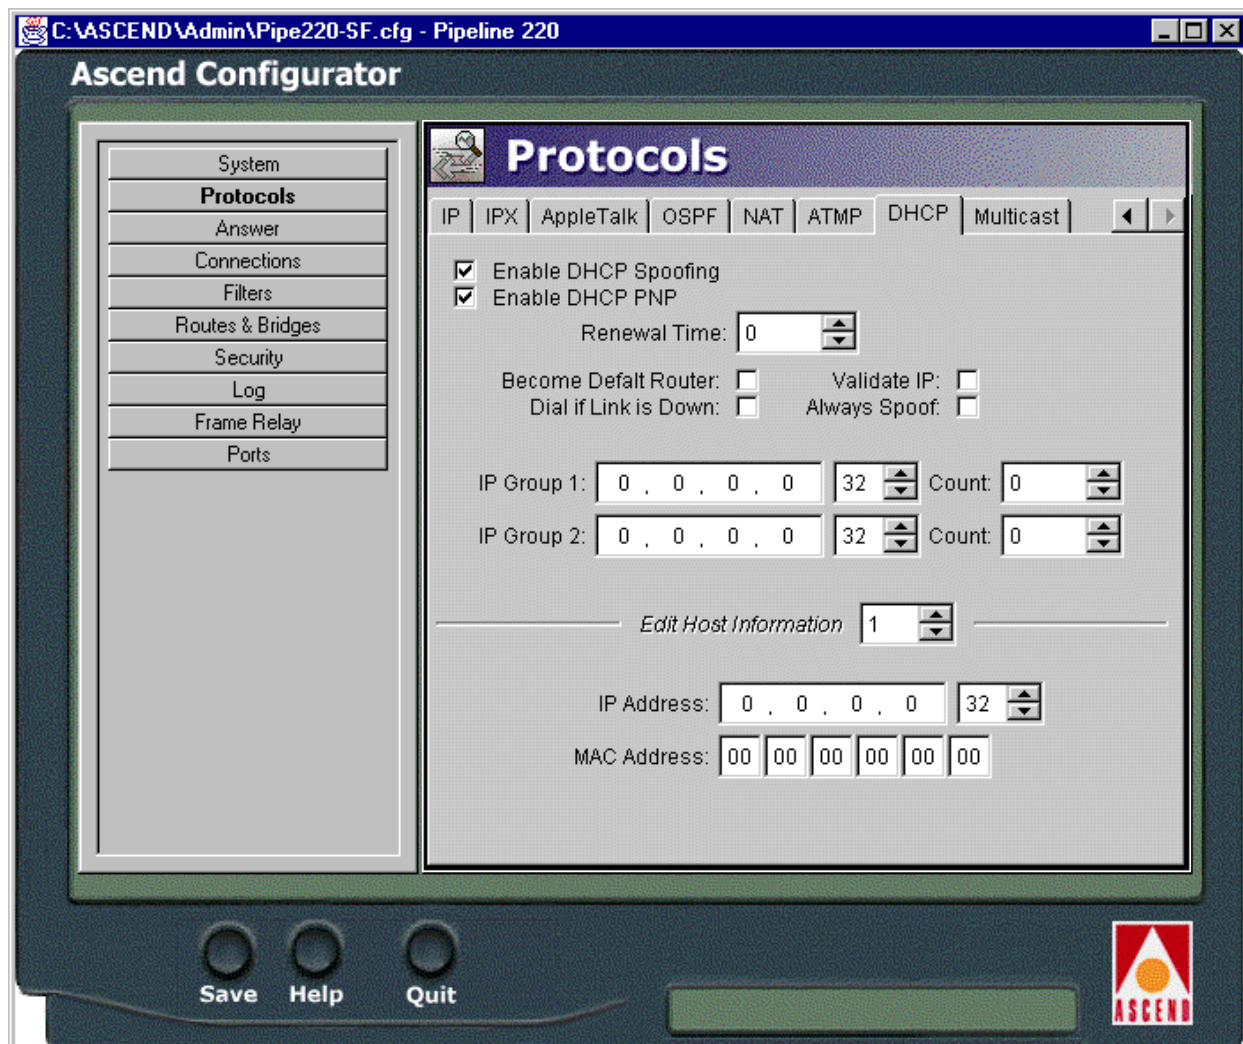
To configure DHCP services, you must first enable it. Then, you configure pools of IP addresses that the Pipeline 220 assigns to any host requesting an address. Additionally, you can configure up to six specific host addresses that receive specific IP addresses.

Enabling DHCP services

To enable DHCP services:

- 1 On the left side of the Configurator, click the Protocols button.
- 2 On the right side of the Configurator, click the IP tab.
- 3 In the lower-right section of the Configurator, click the DHCP button.

System-wide DHCP options appear in the lower-right section of the Configurator:



- 4 Select the “Enable DHCP Spoofing” check box to enable all DHCP services. If you clear the check box, the Pipeline 220 disables all other settings on this screen.
- 5 To enable Plug-and-Play, select the “Enable DHCP PNP” check box.
- 6 Set Renewal Time to the number of seconds a DHCP-assigned IP address is valid before it needs to be renewed. The time applies to DHCP spoofed addresses and DHCP server replies. If the host renews the address before the time expires, the Pipeline 220 provides the same address.

Note: Plug and Play addresses always expire in 60 seconds.

- 7 Select “Become Default Router” to advertise the address of your Pipeline 220 as the default router for all DHCP request packets.
- 8 Select, or clear the Always Spoof check box.
 - If you select the check box, a DHCP server always supplies an IP address for every request, until all IP addresses are exhausted.
 - If you clear the check box, 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.

- 9 If you want the Pipeline 220 to validate that any address it assigns is not in use by another host, check the Validate IP check box.
- 10 Enable the new settings, as described in “Saving your settings” on page 7-2.

Configuring IP address pools

To configure pools of addresses from which the Pipeline 220 assigns IP addresses to requesting hosts:

- 1 Set the IP Group 1 parameter to the first address the Pipeline 220 uses for its first DHCP address pool.
- 2 Set Count to the maximum number of addresses that the Pipeline 220 can assign from this pool.
- 3 If appropriate, set the IP Group 2 parameter to the first address the Pipeline 220 uses for its second DHCP address pool.
- 4 Set Count to the maximum number of addresses that the Pipeline 220 can assign from this pool.
- 5 Enable the new settings, as described in “Saving your settings” on page 7-2.

Assigning specific addresses to particular hosts

To configure specific IP addresses for use by particular hosts:

- 1 Set Edit Host Information to 1.
- 2 Specify the IP address that the Pipeline 220 assigns to the requesting host.
- 3 Specify the Ethernet (MAC) address of the host that is assigned the configured IP address.
- 4 For additional hosts, repeat step 1 through step 3, incrementing the value of the Edit Host Information field. You can specify up to six hosts.
- 5 Enable the new settings, as described in “Saving your settings” on page 7-2.

Local DNS host address table

You can create a local DNS table that provides 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.

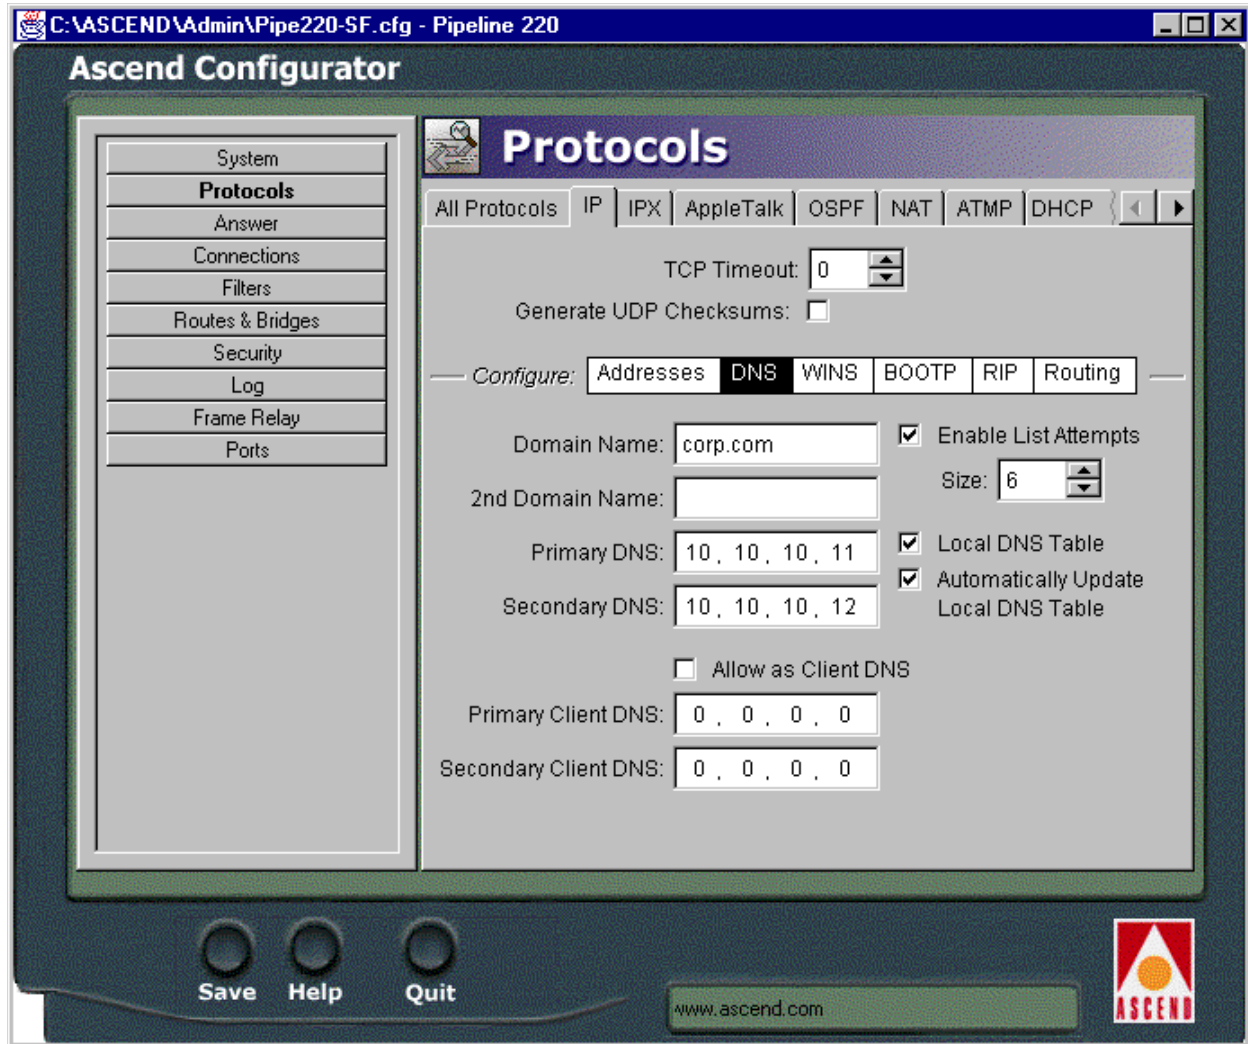
You create the DNS table, from the terminal server, by entering the host names and their IP addresses in the table. A table can contain up to eight entries, with a maximum of 35 IP addresses for each entry. You enter only the first IP address. Any other IP addresses in the list are automatically added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

Also, you can specify that the local DNS table be 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.

To enable and configure the local DNS table:

- 1 On the left side of the Configurator, click the Protocols button.
- 2 On the right side of the Configurator, click the IP tab.
- 3 On the lower-right section of the Configurator, click the DNS button.

System-wide DNS parameters appear in the lower-right section of the Configurator:



- 4 If you select “Enable List Attempts,” the Pipeline 220 displays a list of the DNS IP addresses that appear when you enter the terminal-server command `Dnstab Entry`. For more information, see Chapter 16, “VT100 Interface System Administration.”
- 5 Set Size to the number of entries the Pipeline 220 displays when you enter the terminal-server command `Dnstab Entry`. Enter a number from 1 to 35.
If the DNS 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 server returns 3 entries, the Pipeline 220 clears the local DNS table and enters the three returned addresses 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 Pipeline 220 enters the first four IP addresses into its DNS table for the entry and clears the remaining addresses.
- If you set List Size=1, the list can contain only one IP address. The Pipeline 220 ignores any others returned by the remote DNS. If you reduce the value of the Size parameter to one, only the first IP address is retained. The Pipeline 220 clears all others the next time the it updates the table entry.

6 Select “Local DNS Table.”

7 Select “Automatically Update Local DNS Table.”

When you enable automatic updating, the Pipeline 220 replaces the list of IP addresses for each entry with a list from the remote DNS when the remote DNS successfully resolves a connection to a host named on the table.

8 Enable the new settings, as described in “Saving your settings” on page 7-2.

For additional information, see “DNSTab command” on page 16-40.

User-definable TCP connection retry timeout

Set the Protocols button > IP tab > TCP Timeout parameter to the maximum length of time the Pipeline 220 waits to complete a connection before trying the next address supplied by a DNS server. If the Pipeline 220 cannot connect to the first host in the list, it tries the next (if available), until it connects or times out.

You should set the TCP Timeout parameter on the basis of 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 220 and are lightly loaded, then a short timeout (a few seconds) might be acceptable, 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
- Remote host with characteristics that are not well known.

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

The default value, zero, specifies that the Pipeline 220 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)

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 guarantees 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 220 can be configured to use Network Address Translation (NAT). NAT works as follows:

- When the local host sends packets to the remote network, the Pipeline 220 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 220 translates the official address on the remote network to the host's private address on the local network.

NAT and port routing

A Pipeline 220 performs NAT in the following 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 Frame Relay connection to the Pipeline 220.
- 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.

When using NAT, the Pipeline 220 is the only local host that is visible to the remote network.

How address translation works

When using NAT, the Pipeline 220 replaces, with the official address, the local IP address of any packet received from the Ethernet interface. The Pipeline 220 makes an entry in its internal translation table, and uses the table to keep track of all active NAT sessions.

When the Pipeline 220 receives packets from the WAN, they all have the Pipeline 220's address as their destination. You can configure the Pipeline 220 to route the packets to up to 10 different TCP or UDP ports on specific local servers. You must configure the list of local servers and the UDP and TCP ports each handles. When you configure the Pipeline 220 to route incoming packets, destined for a particular TCP or UDP port, to a specific local server, multiple hosts on the remote network can connect to the server at the same time.

For example, if you configure the Pipeline 220 to route all incoming packets for TCP port 80 (the standard port for HTTP) to port 80 of your local World Wide Web server.

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 received from the WAN and destined for TCP port 119, the well known port for Network News Transfer Protocol, to port 1119 on your local Usenet News server.

You can also define a local default server that handles UDP and TCP ports not listed. If you don't specify any routed ports but do specify a default server, the default server receives all packets that the remote network sends to the Pipeline 220.

The number of simultaneous connections is limited to the size of the translation table.

Translation table size

The Pipeline 220 maintains an internal NAT 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.

Translation-table entries are freed on the basis of the following time outs:

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

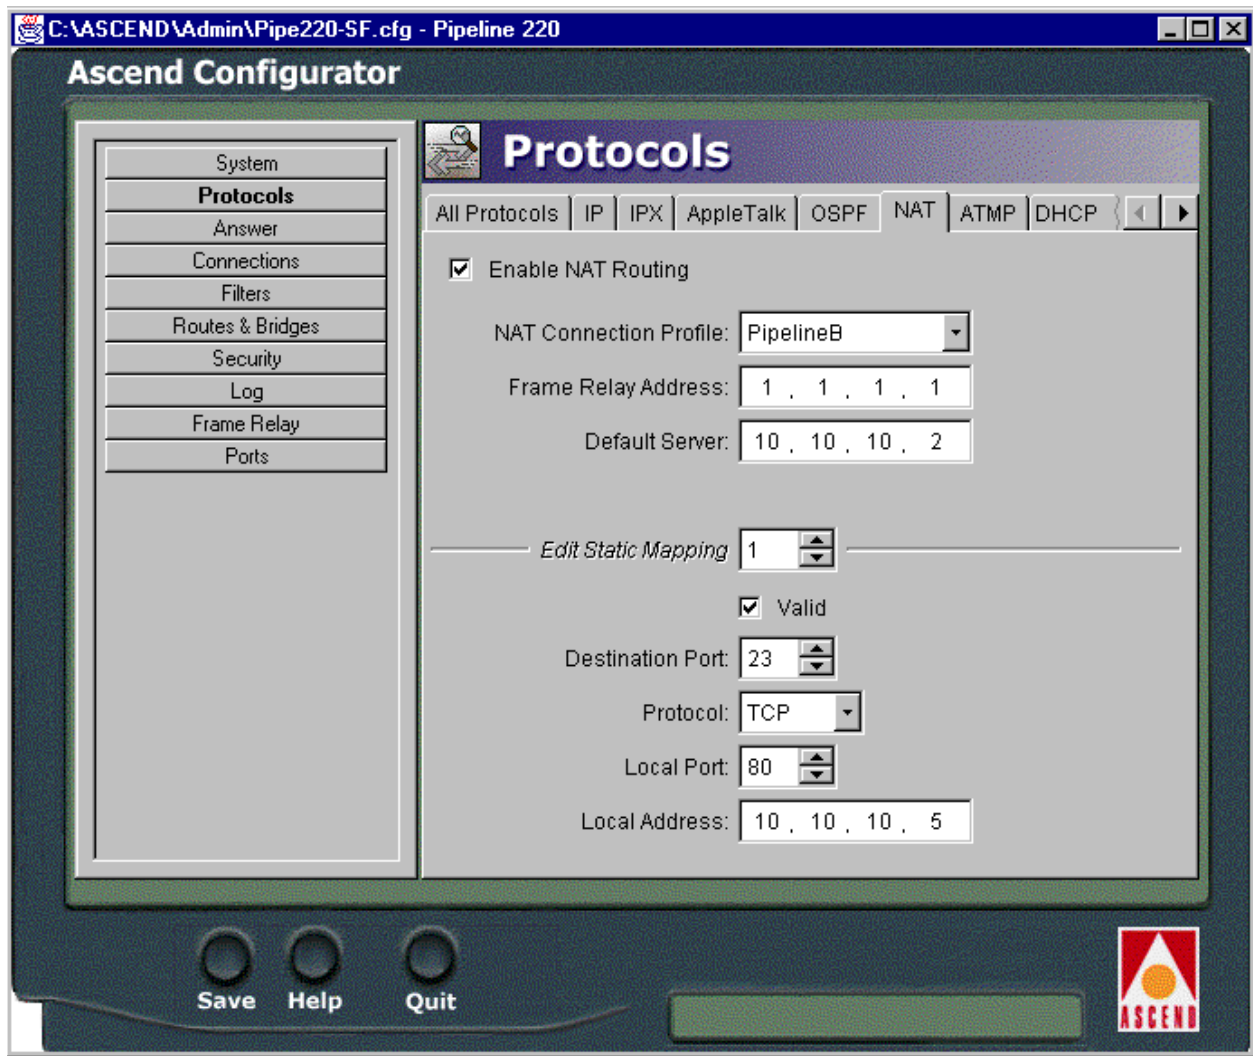
The Pipeline 220 reuses the translation table entries as long as it receives packets that match an entry. All the entries expire when a connection disconnects. Nailed connections are designed not to disconnect.

Configuring NAT

To configure NAT on the Pipeline 220:

- 1** On the left side of the Configurator, click the Protocols button.
- 2** On the right side of the Configurator, click the NAT tab.

System-wide NAT parameters appear on the right side of the Configurator:



- 3 Select "Enable NAT Routing".
- 4 Set NAT Connection Profile to the name of a Connection profile the Pipeline 220 uses to connect to the Network Access Server (NAS).
- 5 Set Frame Relay Address to the official IP address that the Pipeline 220 uses to communicate with the remote network.

Configuring NAT port routing

For remote hosts wishing to access services, you can configure a Pipeline 220 to perform NAT on the local network in one of the following ways:

- Route incoming packets from a remote network for up to 10 different TCP or UDP ports to specific servers and ports on the local network and, optionally, route any remaining packets to a default server.

- Route incoming packets from a remote network, for up to 10 different TCP or UDP ports, to specific servers and ports on the local network.

If the Pipeline 220 receives a packet of the configured type for the configured port, it forwards it to the local IP address and port you specify.

You can configure up to ten static mappings, directing the Pipeline 220 where to forward packets addressed to specific ports. To configure static mappings:

- 1 Set “Edit Static Mapping” to 1. If you have previously configured static mappings, set “Edit Static Mapping” to the next available value.
- 2 Set Destination Port.
Any packet, received from the WAN, that has been sent to the port specified in the Destination Port parameter is forwarded to the port specified in the Local Port parameter.
- 3 Set Protocol to the type of packets to which the Pipeline 220 should apply this mapping.
- 4 Set Local Port to the port to which the Pipeline 220 forwards packets on the local network.
- 5 Set Local Address to the IP address of the host to which the Pipeline 220 forwards packets on the local network.
- 6 Optionally, set Default Server to the IP address of a local server to which the Pipeline 220 routes incoming packets that are *not* routed to a specific server and port.

Note: If you have additional routers on your local area network, you should select the Protocol button > IP tab > RIP button > Ignore Default Route check box. Then, default routes from the ISP will not overwrite the configured NAT route.

- 7 Enable the new settings, as described in “Saving your settings” on page 7-2.

Well-known ports

TCP and UDP ports numbered 0-1023 are assigned by the Internet Assigned Numbers Authority (IANA). In most 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 are registered with the IANA) via FTP from:

`ftp://ftp.isi.edu/in-notes/iana/assignments/port-numbers`

Configuring OSPF Routing

This chapter covers the following topics:

Introduction to OSPF	8-1
Configuring OSPF routing in the Pipeline 220	8-10

Introduction to OSPF

Open Shortest Path First (OSPF) is a second generation Internet routing protocol. The *Open* in its name refers to the fact that OSPF was developed in the public domain as an open specification. *Shortest Path First* refers to an algorithm developed by Dijkstra in 1978 for building a self-rooted shortest-path tree from which routing tables can be derived. (The algorithm is described in “The link-state routing algorithm” on page 8-7.)

RIP limitations solved by OSPF

The rapid growth of the Internet has pushed the Routing Information Protocol (RIP) beyond its capabilities, especially because of the following problems:

Problem	Description and solution
Distance-vector metrics	<p>RIP is a distance-vector protocol, which uses a hop count to select the shortest route to a destination network. RIP always uses the lowest hop count, regardless of the speed or reliability of a link.</p> <p>OSPF is a link-state protocol, which means that OSPF can take into account a variety of link conditions, such as the reliability or speed of the link, when determining the best path to a destination network.</p>
15-hop limitation	<p>A destination that requires more than 15 consecutive hops is considered unreachable, which inhibits the maximum size of a network.</p> <p>OSPF has no hop limitation. You can add as many routers to a network as you want.</p>

Problem	Description and solution
Excessive routing traffic and slow convergence	<p>RIP creates a routing table and then propagates it throughout the internet of routers, hop by hop. The time it takes for all routers to receive information about a topology change is called <i>convergence</i>. A slow convergence can result in routing loops and errors.</p> <p>A RIP router broadcasts its entire routing table every 30 seconds. On a 15-hop network, convergence can be as high as 7.5 minutes. In addition, a large table can require multiple broadcasts for each update, which consumes a lot of bandwidth.</p> <p>OSPF uses a topological database of the network and propagates only changes to the database (as described in “Exchange of routing information” on page 8-4).</p>

Ascend implementation of OSPF

The primary goal for OSPF implementation in this release is to allow the Pipeline 220 to communicate with other routers within a single Autonomous System (AS).

The Pipeline 220 acts as an OSPF internal router with limited border-router capability. For this release, you should not configure the Pipeline 220 as an Area Border Router (ABR), so the Ethernet interface and all of the Pipeline 220 WAN links should be configured in the same area.

The Pipeline 220 does not function as a full AS Border Router (ASBR) at this release. However, it performs ASBR calculations for external routes such as WAN links that do not support OSPF. The Pipeline 220 imports external routes into its OSPF database and flags them as Autonomous System External (ASE). It redistributes those routes via OSPF ASE advertisements, and propagates its OSPF routes to remote WAN routers running RIP.

The Pipeline 220 supports null and simple password authentication.

OSPF features

This section provides a brief overview of OSPF routing to help you configure the Pipeline 220 properly. For full details about how OSPF works, see RFC 1583, “OSPF Version 2,” 03/23/1994, J. Moy.

An Autonomous System (AS) is a group of OSPF routers exchanging information, typically under the control of one company. An AS can include a large number of networks, all of which are assigned the same AS number. All information exchanged within the AS is *interior*.

Exterior protocols are used to exchange routing information between autonomous systems. They are referred to as Exterior Gateway Protocol (EGP). The AS number can be used by border routers to filter out certain EGP routing information. OSPF can make use of EGP data generated by other border routers, and added into the OSPF system as ASEs, as well as static routes configured in the Pipeline 220 or RADIUS.

Security

All OSPF protocol exchanges are authenticated. This means that only trusted routers can participate in the AS's routing. A variety of authentication schemes can be used. In fact, different authentication types can be configured for each area. In addition, authentication provides added security for the routers that are on the network. Routers that do not have the password will not be able to gain access to the routing information, because authentication failure prevents a router from forming adjacencies.

Support for variable length subnet masks

OSPF enables the flexible configuration of IP subnets. Each route distributed by OSPF has a destination and mask. Two different subnets of the same IP network number may have different sizes (different masks). This is commonly referred to as Variable-Length Subnet Masks (VLSM), or Classless Inter-Domain Routing (CIDR). A packet is routed to the best (longest, or most specific) match. Host routes are considered to be subnets whose masks are *all ones* (0xFFFFFFFF).

Note: Although OSPF is very useful for networks that use VLSM, you should assign subnets as contiguously as possible, to prevent excessive link-state calculations by all OSPF routers on the network.

Interior Gateway Protocol (IGP)

OSPF keeps all AS-internal routing information within that AS. All information exchanged within the AS is *interior*.

An AS Border Router (ASBR) is required to communicate with other autonomous systems. To do so, it must use an External Gateway Protocol (EGP), as shown in Figure 8-1. An EGP acts as a shuttle service between autonomous systems.

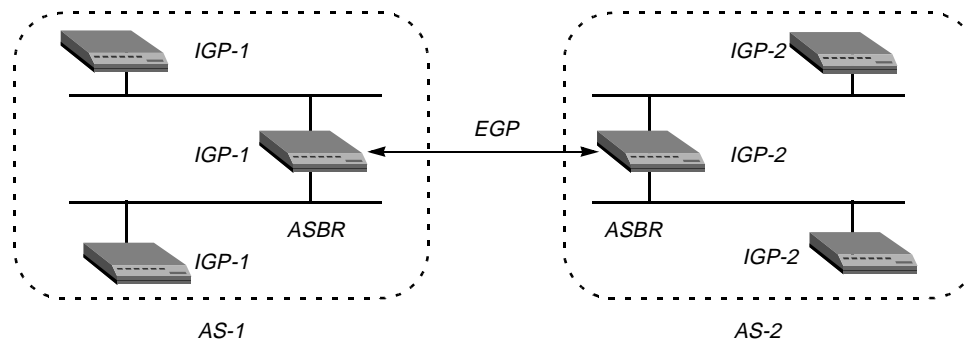


Figure 8-1. Autonomous system border routers

ASBRs perform calculations related to external routes. The Pipeline 220 imports external routes from RIP (for example, when it establishes a WAN link with a caller that does not support OSPF), and always performs the ASBR calculations.

If you must prevent the Pipeline 220 from performing ASBR calculations, you can disable them by clearing the Protocols button > OSPF tab > Global Options button > Enable ASBR check box.

Exchange of routing information

OSPF uses a topological database of the network and propagates only changes to the database. Part of the SPF algorithm involves acquiring neighbors, and then forming an adjacency with one neighbor, as shown in Figure 8-2.




Figure 8-2. Adjacency between neighboring routers

An OSPF router dynamically detects its neighboring routers by sending its Hello packets to the multicast address All OSPF Routers. It then attempts to form adjacencies with some of its newly acquired neighbors.

Adjacency is a relationship formed between selected neighboring routers for the purpose of exchanging routing information. Not every pair of neighboring routers become adjacent. Adjacencies are established during network initialization, in pairs, between two neighbors. As the adjacency is established, the neighbors exchange databases and build a consistent, synchronized database between them.

When an OSPF router detects a change on one of its interfaces, it modifies its topological database and multicasts the change to its adjacent neighbor, which in turn propagates the change to its adjacent neighbor until all routers within an area have synchronized topological databases. The result is quick convergence among routers. OSPF routes can also be summarized in Link-State Advertisements (LSAs).

Designated and backup designated routers

In OSPF terminology, a broadcast network is any network that has more than two OSPF routers attached and supports the capability to address a single physical message to all of the attached routers.




Figure 8-3. Designated and backup designated routers

The Pipeline 220 can function as a Designated Router (DR) or Backup Designated Router (BDR). However, many sites choose to assign a LAN-based router for these roles in order to dedicate the Pipeline 220 to WAN processing.

To reduce the number of adjacencies each router must form, OSPF calls one of the routers the designated router. A designated router is elected as routers are forming adjacencies, and then all other routers establish adjacencies only with the designated router. This simplifies the routing table update procedure and reduces the number of link-state records in the database. The designated router plays other important roles in addition to reducing the overhead of OSPF link-state procedures. For example, other routers send link-state advertisements to only the designated router by using the *all-designated-routers* multicast address of 224.0.0.6.

To prevent the designated router from becoming a serious liability to the network if it fails, OSPF elects a backup designated router at the same time it elects the designated router. Other routers maintain adjacencies with both the designated router and its backup router, but the backup router leaves as many of the processing tasks as possible to the designated router. If the designated router fails, the backup immediately becomes the designated router and a new backup is elected.

The administrator chooses which router is to be the designated router on the basis of the processing power, speed, and memory of the system, and then assigns priorities to other routers on the network in case the backup designated router is also down at the same time.

Configurable metrics

The administrator assigns a cost to the output side of each router interface. The lower the cost, the more likely is the interface to be used to forward data traffic. Costs can also be associated with the externally derived routing data.

The OSPF cost can also be used for preferred-path selection. If two paths to a destination have equal costs, you can assign a higher cost to one of the paths to configure it as a backup to be used only when the primary path is not available.

Figure 8-4 shows how costs are used to direct traffic over high-speed links. For example, if Router-2 in Figure 8-4 receives packets destined for Host B, it routes them through Router-1 across two T1 links (Cost=20), rather than across one 56Kbps B-channel to Router-3 (Cost=240).

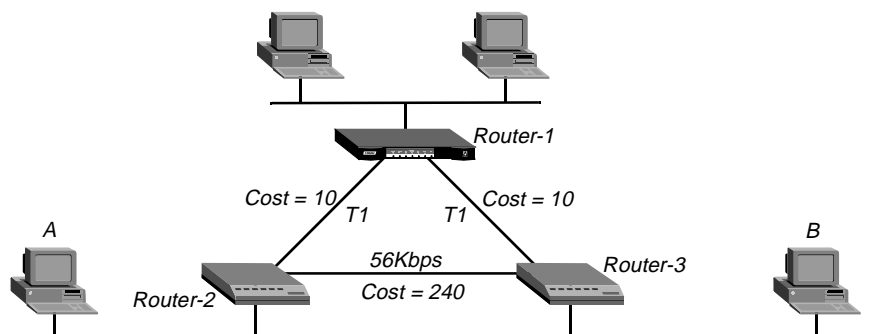


Figure 8-4. OSPF costs for different types of links

The Pipeline 220 has a default cost of 1 for a connected route (Ethernet) and 10 for a WAN link. You might want to reflect the bandwidth of a connection in its cost assignment. For example, for a single B-channel connection, the cost would be 24 times greater than for a T1 link.

Note: Be careful when assigning costs. Incorrect cost metrics can cause delays and congestion on the network.

Hierarchical routing (areas)

If a network is large, the size of the database, time required for route computation, and related network traffic can become excessive. An administrator can partition an AS into areas to provide hierarchical routing connected by a backbone. (But do not use the Pipeline 220 to connect to the backbone.)

The backbone area is special and always has the area number 0.0.0.0. Other areas are assigned area numbers that are unique within the autonomous system.

Each area acts like its own network. That is, all area-specific routing information stays within the area, and all routers within an area must have a synchronized topological database. To tie the areas together, some routers belong to the backbone area as well as to their own respective areas. These routers are Area Border Routers (ABRs). In Figure 8-5, all of the routers are ABRs. If the ABRs and area boundaries are set up correctly, link-state databases are unique to an area.




Figure 8-5. Dividing an AS into areas

Note: With this release, you should not configure the Pipeline 220 as an ABR. You should use the same area number for the Ethernet interface of the Pipeline 220 and each of its WAN links. That number does not have to be the backbone. The Pipeline 220 can reside in any OSPF area.

Stub areas

To reduce the cost of routing, OSPF supports stub areas, in which a default route summarizes all external routes. For areas that are connected to the backbone by only one ABR (that is, the area has one exit point), there is no need to maintain information about external routes. Stub areas are similar to regular areas except that the routers do not enter external routes in the area's databases.

To prevent flooding of external routes throughout the AS, you can configure an area as a stub when there is a single exit point from the area, or when the choice of exit point need not be made on a per-external-destination basis. You might need to specify a stub area with no default cost (StubNoDefault) if the area has more than one exit point.

In a stub area, routing to AS-external destinations is based on a per-area default cost. The per-area default cost is advertised to all routers within the stub area by a border router, and is used for all external destinations.

Note: If the Pipeline 220 supports external routes across its WAN links, you should not configure it in a stub area. Because an ABR configuration is not currently recommended for the Pipeline 220, the area in which it resides should not be a stub area if any of its links are AS-external.

The link-state routing algorithm

Link-state routing algorithms require that all routers within a domain maintain synchronized (identical) topological databases, and that the databases describe the complete topology of the domain. An OSPF router's domain may be an AS or an area within an AS.

OSPF routers exchange routing information and build Link-state databases. Link-state databases are synchronized between pairs of adjacent routers (as described in “Exchange of routing information” on page 8-4). In addition, each OSPF router uses its link-state database to calculate a self-rooted tree of shortest paths to all destinations, as shown in Figure 8-6.

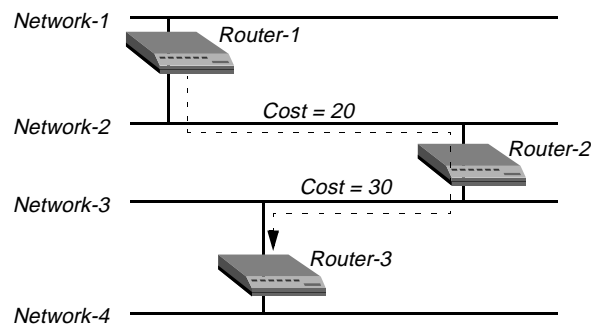


Figure 8-6. Sample network topology

The routers then use the trees to build their routing tables, as shown in Table 8-1.

Table 8-1. Link state databases for network topology in Figure 8-6

Router-1	Router-2	Router-3
Network-1/Cost 0	Network-2/Cost0	Network-3/Cost 0
Network-2/Cost 0	Network-3/Cost0	Network-4/Cost 0
Router-2/Cost 20	Router-1/Cost 20	Router-2/Cost 30
	Router-3/Cost 30	

Table 8-2, Table 8-3, and Table 8-4 show another example of self-rooted shortest-path trees calculated from link-state databases, and the resulting routing tables. Actual routing tables also contain externally derived routing data, which is advertised throughout the AS but kept separate from the Link-state data. Also, each external route can be tagged by the advertising router, enabling the passing of additional information between routers on the boundary of the AS.

Table 8-2. Shortest-path tree and resulting routing table for Router-1

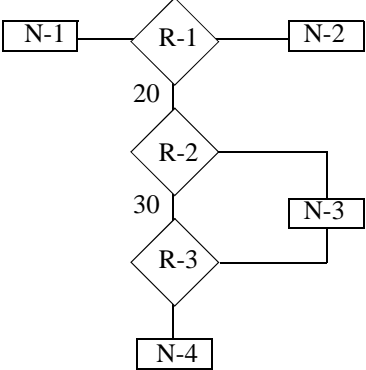
	Destination	Next Hop	Metric
	Network-1	Direct	0
	Network-2	Direct	0
	Network-3	Router-2	20
	Network-4	Router-2	50

Table 8-3. Shortest-path tree and resulting routing table for Router-2

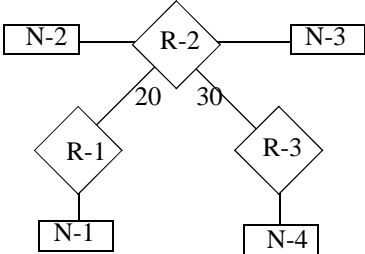
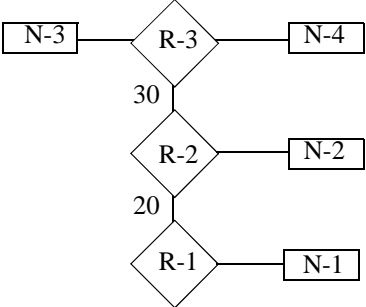
	Destination	Next Hop	Metric
	Network-1	Router-1	20
	Network-2	Direct	0
	Network-3	Direct	0
	Network-4	Router-2	30

Table 8-4. Shortest-path tree and resulting routing table for Router-3

	Destination	Next Hop	Metric
	Network-1	Router-2	50
	Network-2	Router-2	30
	Network-3	Direct	0
	Network-4	Direct	0

Configuring OSPF routing in the Pipeline 220

This section provides brief descriptions of the OSPF routing parameters, and shows how to use the parameters in sample configurations. For additional information about each parameter, see the *Configurator Online Help*.

Understanding the OSPF routing parameters

This section provides some background information about the OSPF parameters. You can access the parameters either through the Protocols button > OSPF tab or through the Connections button > OSPF tab. The parameters are the same, but some of the default values are different. For OSPF routing, you configure the following settings:

Setting	Description
Enabling OSPF on an interface	OSPF is turned off by default. To enable it on an interface, set RunOSPF to Yes.
Specifying an area number and type	Area sets the area ID for the interface. The format for this ID is dotted decimal, but it is not an IP address. (For a description of areas, see “Hierarchical routing (areas)” on page 8-6.) AreaType specifies the type of area: Normal, Stub, or StubNoDefault. (For descriptions, see “Stub areas” on page 8-6.)
Intervals for communicating with an adjacent router	HelloInterval specifies how frequently, in seconds, the Pipeline 220 sends out Hello packets on the specified interface. OSPF routers use Hello packets to dynamically detect neighboring routers in order to form adjacencies. DeadInterval specifies how many seconds the Pipeline 220 will wait before declaring its neighboring routers down after it stops receiving their Hello packets (For background information, see “Exchange of routing information” on page 8-4.)
Priority	The routers in the network use the Priority value to elect a Designated Router (DR) and Backup Designated Router (BDR). Assigning a priority of 1 would place the Pipeline 220 near the top of the list of possible designated routers. (Currently, you should assign a larger number.) Acting as a DR or BDR significantly increases the amount of OSPF overhead for the router. (For a discussion of the functions of DRs and BDRs, see “Designated and backup designated routers” on page 8-4.)
Authentication type and key	You can specify that the Pipeline 220 supports OSPF router authentication, and the key it will look for in packets to support that authentication. See “Security” on page 8-3.

Setting	Description
Cost of the route on this interface	This parameter specifies the link-state or output cost of a route. Assign realistic costs for each interface that supports OSPF. The lower the cost, the higher the likelihood of using that route to forward traffic. See “Configurable metrics” on page 8-5.
Autonomous System External route (ASE) type and tag	<p>ASEs are used only when OSPF is turned off on a particular interface. When OSPF is enabled, the ASE parameters are not applicable.</p> <p>ASE-type specifies the type of metric that the Pipeline 220 advertises for external routes. A Type 1 external metric is expressed in the same units as the link-state metric (the same units as interface cost). A Type 2 external metric is considered larger than any link state path. Use of Type 2 external metrics assumes that routing between autonomous systems is the major cost of routing a packet, and eliminates the need for conversion of external costs to internal link-state metrics.</p> <p>ASE-tag is a hexadecimal number used to tag external routes for filtering by other routers.</p>
Transit delay	Specify the estimated number of seconds it takes to transmit a Link State Update Packet over this interface, taking into account transmission and propagation delays. On a connected route, you can leave the default of 1.
Retransmit interval	Specify the number of seconds between retransmissions of Link-State Advertisements, Database Description, and Link State Request Packets.
OSPF global option for disabling ASBR calculations	<p>Autonomous System Border Routers (ASBRs) perform calculations related to external routes. The Pipeline 220 imports external routes from RIP (for example, when it establishes a WAN link with a caller that does not support OSPF) and the ASBR calculations are always performed. If you must prevent the Pipeline 220 from performing ASBR calculations, disable them by clearing the Protocols button > OSPF tab > Global Options button > Enable ASBR check box.</p>

Examples of adding the Pipeline 220 to an OSPF network

This section shows how to add a Pipeline 220 to your OSPF network. It assumes that you know how to configure the Pipeline 220 with an appropriate IP address (as described in Chapter 6, “Configuring IP Routing”). The procedures in this section are examples based on Figure 8-7. Configuring the unit labeled Pipeline 220-1 in Figure 8-7. To apply one or more of the procedures to your network, enter the appropriate settings instead of the ones shown.




Figure 8-7. A sample OSPF setup

In Figure 8-7, all OSPF routers are in the same area (the backbone area), so the units will all form adjacencies and synchronize their databases together.

Note: All OSPF routers in Figure 8-7 have RIP turned off. OSPF can learn routes from RIP without the added overhead of running RIP.

Configuring OSPF on the Ethernet interface

The Pipeline 220 Ethernet interface in Figure 8-7 is in the OSPF backbone area. Although the RFC states no limitation regarding the number of routers in the backbone area, you should keep the number relatively small, because changes that occur in area zero propagate throughout the AS.

Another way to configure the same units would be to create a second area (such as 0.0.0.1) for one of the existing OSPF routers, and add the Pipeline 220 to that area. You could then assign the same area number (0.0.0.1) to all OSPF routers reached through the Pipeline 220 across a WAN link.

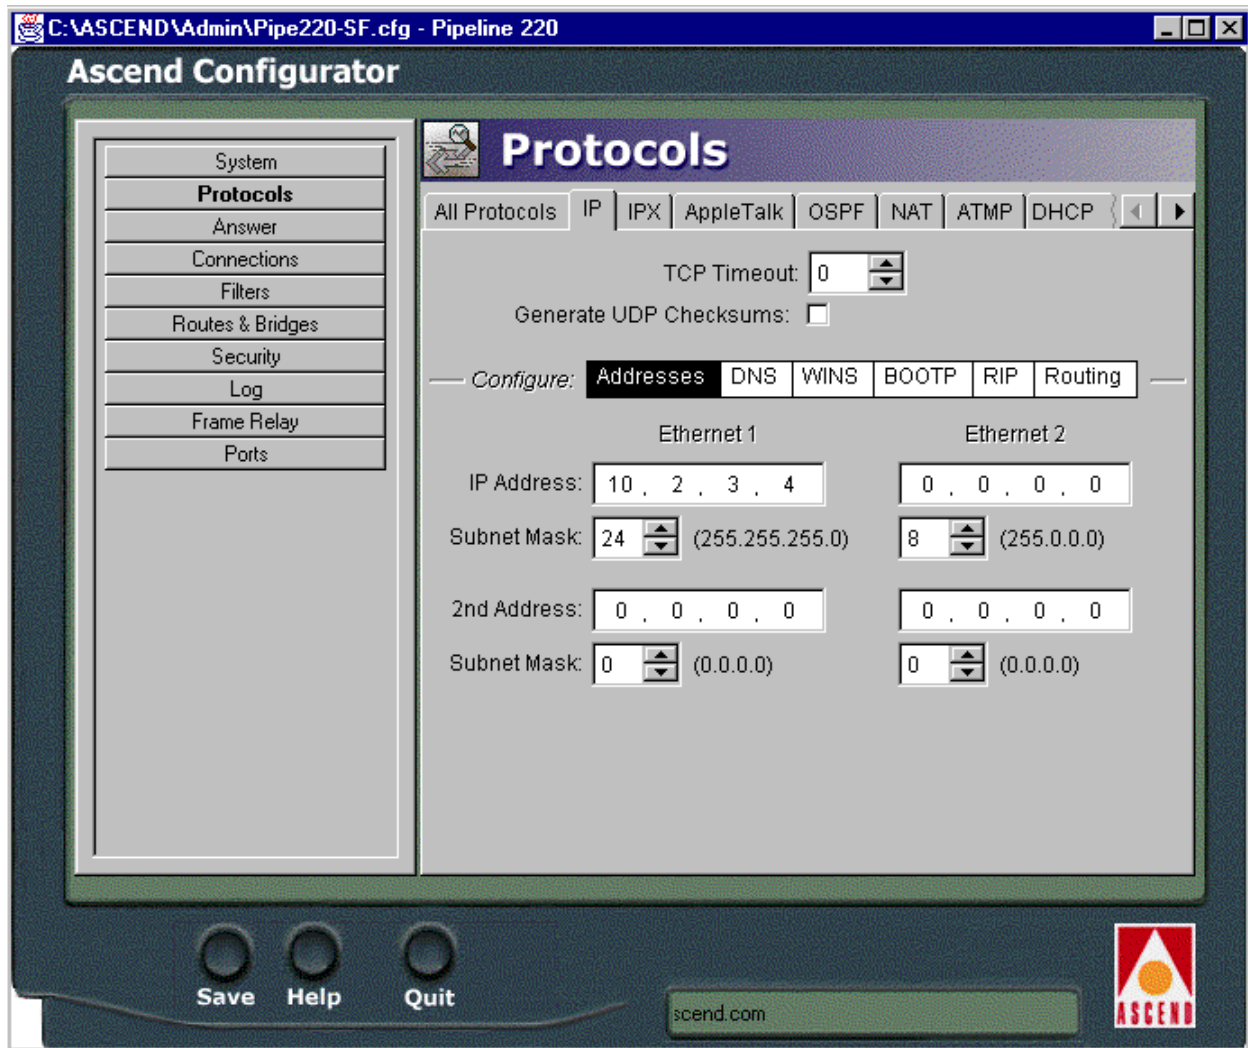
After you configure the Pipeline 220 as an IP host on that interface, you can configure it, in the Ethernet profile, as an OSPF router in the backbone area. To configure the Pipeline 220 as an OSPF router on Ethernet, first verify that you have configured the Pipeline 220 as an IP router. Then enable OSPF routing, configure authentication, configure OSPF cost and intervals, and save your settings.

Verifying IP connectivity

To verify IP connectivity:

- 1 On the left side of the Configurator, click the Protocols button.
- 2 On the right side of the Configurator, click the IP tab.

- 3 In the lower-right section of the Configurator, click the Addresses button.
IP-address parameters appear in the lower-right section of the Configurator:



Each working Ethernet interface must be configured with a valid IP address and subnet mask. To validate the IP addresses, use another IP host on your network. Ping or Telnet to each configured Pipeline 220 address. If you receive no response from the Pipeline 220, and believe the addresses are valid for your environment, see your network administrator.

Note: It is not necessary to enable RIP and OSPF simultaneously, and disabling RIP reduces the processing overhead of the Pipeline 220. If you have routers on your network that support RIP only, the Pipeline 220 uses OSPF to learn routes from RIP, incorporating them into its routing table, assigning them an external metric, and tagging them as external routes. (For more information, see Chapter 6, “Configuring IP Routing.”)

Enabling OSPF routing

To enable OSPF routing on the Pipeline 220.

- 1 On the left side of the Configurator, click the Protocols button.
- 2 On the right side of the Configurator, click the OSPF tab.

Configuring OSPF Routing

Configuring OSPF routing in the Pipeline 220

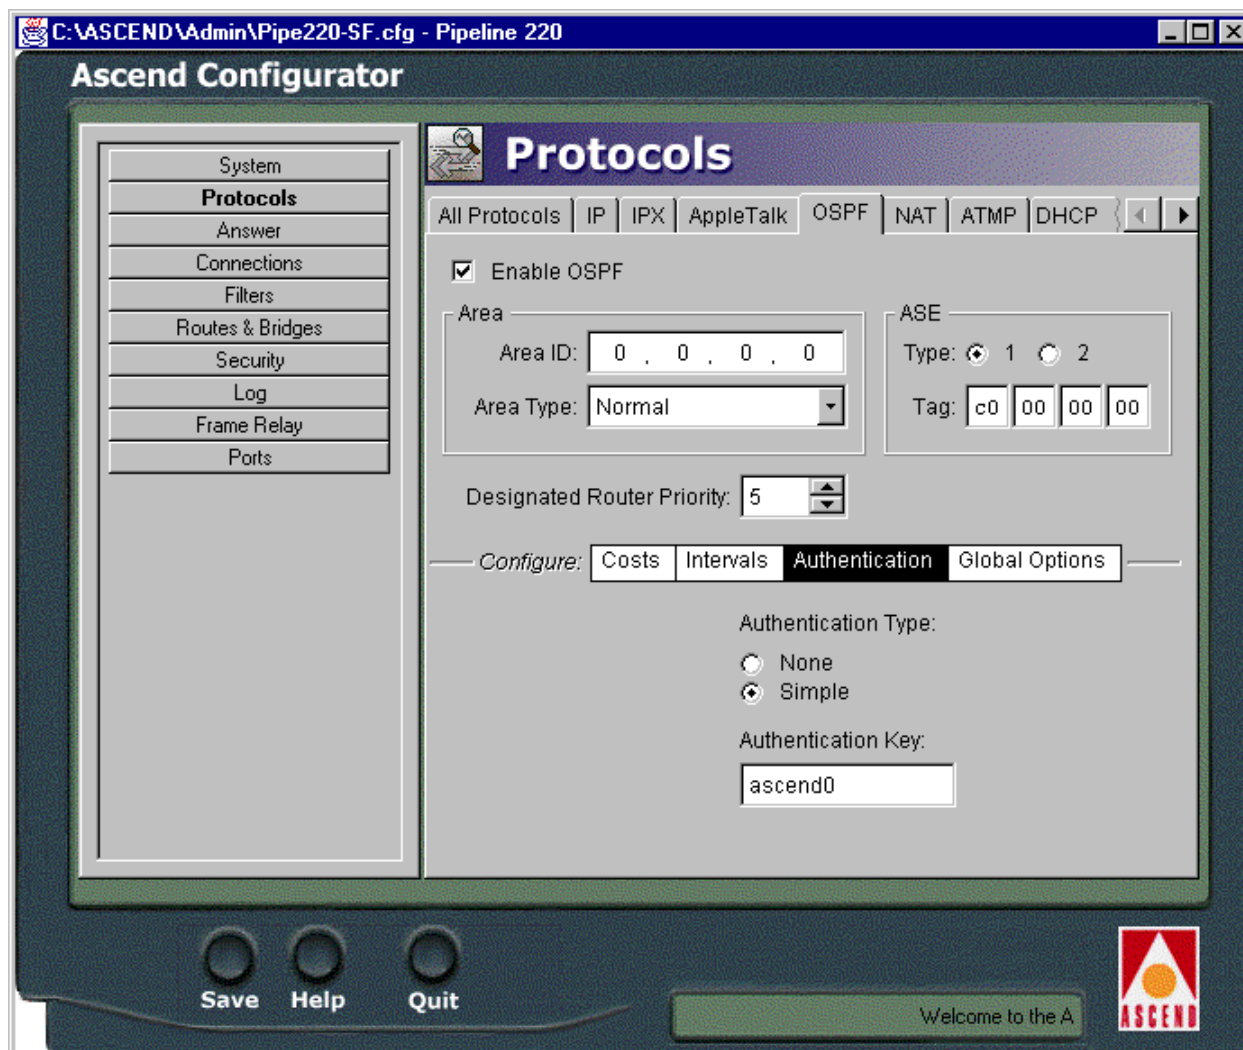
- 3 Select the Enable OSPF check box.
- 4 Specify the area number and area type for the Ethernet interfaces.
For example, when the Ethernet interface is in the backbone area, set Area ID to 0.0.0.0. (Because the backbone area is not a *stub* area, so you should leave the setting at its default value of 0.0.0.0. (For background information, see “Stub areas” on page 8-6.)

Configuring authentication

If access to the backbone area requires authentication, specify the password:

- 1 In the lower-right section of the Configurator, click the Authentication button.

OSPF-authentication parameters appear in the lower-right section of the Configurator:



- 2 Set Authentication Type to Simple.
If authentication is not required, set Authentication Type to None.
- 3 Specify the password under Authentication Key.

Configuring OSPF costs and intervals


Configure metrics to identify characteristics of the route and specify the likelihood that the Pipeline 220 should use a particular route to forward traffic. It is common practice to configure multiple routes to identical networks, and metrics show the Pipeline 220 how to rank each route. To configure the metrics for a route:

- 1 On the lower-right section of the Configurator, click the Costs button.
- 2 Set Output Cost to a number from 1 to 16777214.
An Ethernet-connected route is typically assigned a cost of 1. This is the default.
- 3 On the lower-right section of the Configurator, click the Intervals button.
- 4 Set Transit Delay to the number of seconds it takes to transmit a Link State Update (LSU) Packet over this interface.
This value should take into account transmission and propagation delays. The default is 1.
- 5 Set Retransmit Interval to the number of seconds between retransmissions of OSPF packets.
OSPF uses this value for LSA transmissions and when retransmitting OSPF Database Description packets and Link State Request packets. The default is 5.
- 6 Enable the new settings, as described in “Saving the settings” on page 8-15.
When you upload the changes, the Pipeline 220 comes up as an OSPF router on that interface. It forms adjacencies and builds its routing table.

Saving the settings

The Pipeline 220 settings you configure with the Ascend Configurator do not take effect until you upload them to the Pipeline 220. You use the Save command to upload the changes. Alternatively, you can save the setting to a text file for subsequent use, or to another device. To save the settings:

- 1 In the lower-left corner of the Configurator, click Save
A dialog box appears, prompting you for a save method:



- 2 Select one of the following:

Configuring OSPF Routing

Configuring OSPF routing in the Pipeline 220

- If you are uploading the configuration to the same device, select “Upload changes to <Pipeline 220 name>.”
 - If you are saving the configuration to a text file on your computer, select “Save a copy under a new filename” and enter the filename to which the configuration should be saved.
 - If you are uploading the configuration to a different device, select “Upload to the Ascend product at:” and enter the IP address of the device.
- 3** Click Save.

If you selected “Upload changes to <Pipeline 220 name>” in step 2, the Ascend

Configurator uploads the configuration file to the Pipeline 220 via TFTP. If you selected one of the other options, the Configurator sends the file to the location you specified. If you specified a remote device, transmission is via TFTP.

Configuring OSPF across the WAN

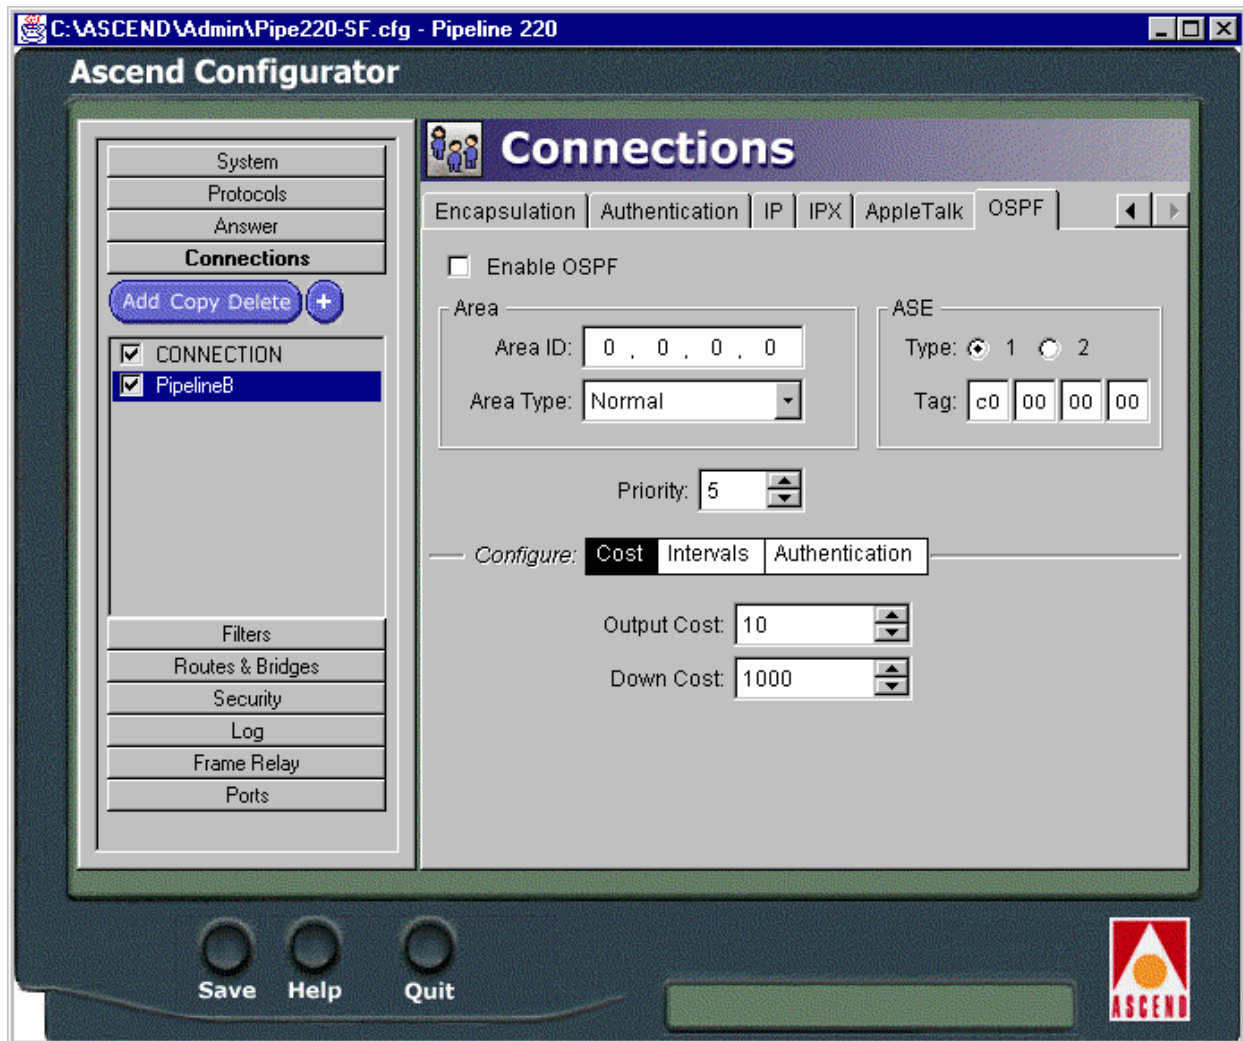
The WAN interface of the Pipeline 220 is a point-to-point network. A point-to-point network is any network that joins a single pair of routers. Such networks typically do not provide a broadcasting or multicasting service, so all advertisements are sent point to point.

An OSPF WAN link has a default Output Cost of 10. You can assign a higher cost to reflect a slower, lower-bandwidth connection or a lower cost to set up a preferred route to a certain destination. If the cost of one route is lower than that of another to the same destination, the higher-cost route will not be used unless route preferences change the equation.

To configure OSPF on a WAN link, you specify settings in a Connection profile. In Figure 8-7, the Pipeline 220 is connecting to another Pipeline 220 unit across a nailed T1 link. To configure the Connection profile:

- 1** On the left side of the Configurator, click the Connections button.
- 2** Highlight an existing Connection profile, or add a new one by clicking the Add button.
- 3** On the right side of the Configurator, click the IP tab.
You must configure IP parameters and OSPF parameters.
- 4** In the lower-right section of the Configurator, click the Addresses button.
- 5** Set IP Address to the address of the remote device’s IP address. For example:
IP Address=10.2.3.100
Subnet Mask=8
(For detailed information, see Chapter 6, “Configuring IP Routing.”)

- 6 On the right side of the Configurator, click the OSPF tab.
OSPF parameters appear in the lower-right section of the Configurator:



- 7 Select the Enable OSPF check box.
- 8 Specify the area number for the remote device and the area type.
The area number must always be specified in dotted-quad format similar to an IP address.
Note: With this release, you should use the same area number for the Ethernet interface of the Pipeline 220 and each of its WAN links.
- 9 If you require authentication for access to the backbone area, click the Authentication button in the lower-right section of the Configurator.
- 10 Specify the password.
For example:
Authentication Type=Simple
Authentication Key=ascend0
If authentication is not required, set Authentication Type=None.
- 11 On the lower-right section of the Configurator, click the Cost button.
- 12 Set Output Cost for the route to the remote device.

You should set the cost of a WAN route to at least 10.

- 13 Enable your changes, as described in “Saving the settings” on page 8-15.

Of course, the remote device must also have a comparable Connection profile to connect to the Pipeline 220.

Configuring a WAN link that doesn't support OSPF

In Figure 8-7, the Pipeline 220 has a Connection profile for connecting to a remote Pipeline unit across a BRI link. The remote Pipeline is an IP router that uses RIP-v2 to transmit routes. The route to this network, as well as any routes the Pipeline 220 learns about from the remote Pipeline, are ASEs (external to the OSPF system).

To enable OSPF to add the RIP-v2 routes to its routing table, configure RIP-v2 normally in this Connection profile. OSPF will import all RIP routes as Type-2 ASEs.

In this example, RIP is turned off on the link and ASE information is configured explicitly.

To, configure the Connection to the remote device:

- 1 On the left side of the Configurator, click the Connections button.
- 2 On the right side of the Configurator, click the IP tab.
You must configure IP parameters and OSPF parameters.
- 3 In the lower-right section of the Configurator, click the Addresses button.
- 4 Set IP Address to the address of the remote device's IP address. For example:
IP Address=10.2.3.4
Subnet Mask=24
For detailed information, see Chapter 6, “Configuring IP Routing.”

Note: The Connections button > OSPF tab includes two ASE parameters that are active only when OSPF is *not* running on a link. After you configure these parameters, the Connection profile route is advertised whenever the Pipeline 220 is up.

- 5 On the right side of the Configurator, click the OSPF tab.
- 6 Clear the Enable OSPF check box.
- 7 On the lower-right section of the Configurator, click the Cost button.
- 8 Set Output Cost for the route to the remote Pipeline.
- 9 In the upper-right section of the Configurator, set ASE Type for this route.
ASE Type specifies the type of metric to be advertised for an external route.
A Type 1 external metric is expressed in the same units as the link state metric (the same units as interface cost). Type 1 is the default.
A Type 2 external metric is considered larger than any link state path. Use of Type 2 external metrics assumes that routing outside the AS is the major cost of routing a packet, and eliminates the need for conversion of external costs to internal link-state metrics.
- 10 Enter an ASE Tag for this route.
The ASE Tag is a hexadecimal number that shows up in management utilities and *flags* this route as external. It can also be used by border routers to filter this record.
- 11 Enable your changes, as described in “Saving the settings” on page 8-15.

Of course, the remote Pipeline 220 unit must also have a comparable Connection profile to connect to the local device.

Setting Up IP Multicast Forwarding

This chapter covers the following topics:

Overview	9-1
Understanding the multicast parameters	9-2
Forwarding from an MBONE router on Ethernet	9-4
Forwarding from an MBONE router on a WAN link	9-8

Overview

The multicast backbone (MBONE) is a virtual network layered on top of the Internet to support IP multicast routing across point-to-point links. It is used for transmitting audio and video on the Internet in real time, because multicasting is a much cheaper and faster way to communicate the same information to multiple hosts.

When using the MBONE, the Pipeline 220 looks like a multicast client. It responds as a client to Internet Group Membership Protocol (IGMP) packets it receives from MBONE routers, which may be IGMP version-1 or version-2, including IGMP MTRACE (multicast trace) packets.

To multicast clients on a WAN or Ethernet interface, the Pipeline 220 looks like a multicast router. Like a router, it sends those clients IGMP queries, receives responses, and forwards multicast traffic. For the current release of the Pipeline 220 software, multicast clients are not allowed to source multicast packets. If they do, the Pipeline 220 discards the packets.

Understanding the multicast parameters

This section provides some background information about multicast parameters.

(For detailed information about each parameter, see the *Configurator Online Help*.)

Enabling multicast forwarding

The Forwarding parameter turns on multicast forwarding in the Pipeline 220.

When you change the Forwarding parameter from No to Yes, the multicast subsystem reads the values in the Ethernet profile and initiates the forwarding function.

Note: If you modify a multicast value in the Ethernet profile, you must set this parameter to No and back to Yes again to force a read of the new value.

Specifying the MBONE interface

The multicast router must connect to the MBONE. It provides its clients with an interface to the MBONE. A Pipeline 220 that resides across the WAN from the multicast router must have a resident Connection profile that defines a connection to the multicast router, and the Mbone Profile parameter must specify that Connection profile. If the Mbone Profile name is null, and Multicast Forwarding is turned on, the Pipeline 220 assumes that its connection to the Ethernet is the MBONE interface.

Monitoring the multicast heartbeat

When it is running as a multicast forwarder, the Pipeline 220 continually receives multicast traffic. The heartbeat-monitoring feature enables the administrator to monitor possible connectivity problems by continuously polling for this traffic and generating an SNMP alarm trap if there is a traffic breakdown. Following is a sample SNMP alarm trap:

```
Trap type: TRAP_ENTERPRISE
Code: TRAP_MULTICAST_TREE_BROKEN (19)
Arguments:
1) Multicast group address being monitored (4 bytes).
2) Source address of last heartbeat packet received (4 bytes).
3) Slot time interval configured in seconds (4 bytes).
4) Number of slots configured (4 bytes).
5) Total number of heartbeat packets received before the Pipeline
220 started sending SNMP Alarms (4bytes).
```

Note: Heartbeat monitoring is optional. It is not required for multicast forwarding.

To set up heartbeat monitoring, you configure several parameters that define which packets will be monitored, how often and for how long to poll for multicast packets, and the threshold for generating an alarm. Following are the parameters you use to specify these settings:

Setting	Parameters
Which packets will be monitored	HeartBeat Address specifies a multicast address. If the parameter is specified, the Pipeline 220 listens for packets to and from this group. HeartBeat UDP port specifies a UDP port number. If it is specified, the Pipeline 220 listens only to packets received through that port. Source Addr and Source Mask specify an IP address and subnet mask. If they are specified, the Pipeline 220 ignores packets from that source for monitoring purposes.
How often and for how long to poll for multicast packets	HeartBeat Slot Time specifies an interval (in seconds). The Pipeline 220 polls for multicast traffic, waits for the duration of the interval, and then polls again. HeartBeat Slot Count specifies how many times to poll before comparing the number of heartbeat packets received to the Alarm Threshold.
The threshold for generating an alarm	Heartbeat Alarm Threshold specifies a number. If the number of monitored packets falls below this number, the SNMP alarm trap is sent.

Configuring multicast forwarding on a client interface

Each local or WAN interface that supports multicast clients must have the Client (or Multicast Client) parameter set to Yes. With this setting, the Pipeline 220 begins handling IGMP requests and responses on the interface. It does not begin forwarding multicast traffic until the rate limit is set.

The Rate Limit parameter specifies the rate at which the Pipeline 220 accepts multicast packets from its clients. It does not affect the MBONE interface. By default, the Rate Limit parameter is set to 100. This disables multicast forwarding on the interface. The forwarder handles IGMP packets, but does not accept packets from clients or forward multicast packets from the MBONE router.

To begin forwarding multicast traffic on the interface, you must set the Rate Limit parameter to a number less than 100. For example, if you set it to 5, the Pipeline 220 accepts a packet from multicast clients on the interface every 5 seconds. Any subsequent packets received in that 5-second window are discarded.

An implicit priority setting for dropping multicast packets

For high-bandwidth data, voice, and audio multicast applications, the Pipeline 220 supports both multicast rate limiting and prioritized packet dropping. If the Pipeline 220 is the receiving device under extremely high loads, it drops packets according to a priority ranking, which is determined by the following UDP port ranges:

- Traffic on ports 0–16384 (unclassified traffic) has the lowest priority (50).

- Traffic on ports 16385–32768 (Audio traffic) has the highest priority (70).
- Traffic on ports 32769–49152 (Whiteboard traffic) has medium priority (60).
- Traffic on ports 49153–65536 (Video traffic) has low priority (55).

Forwarding from an MBONE router on Ethernet

Figure 9-1 shows a local multicast router on one of the Pipeline 220 unit's Ethernet interfaces, and multicast clients.

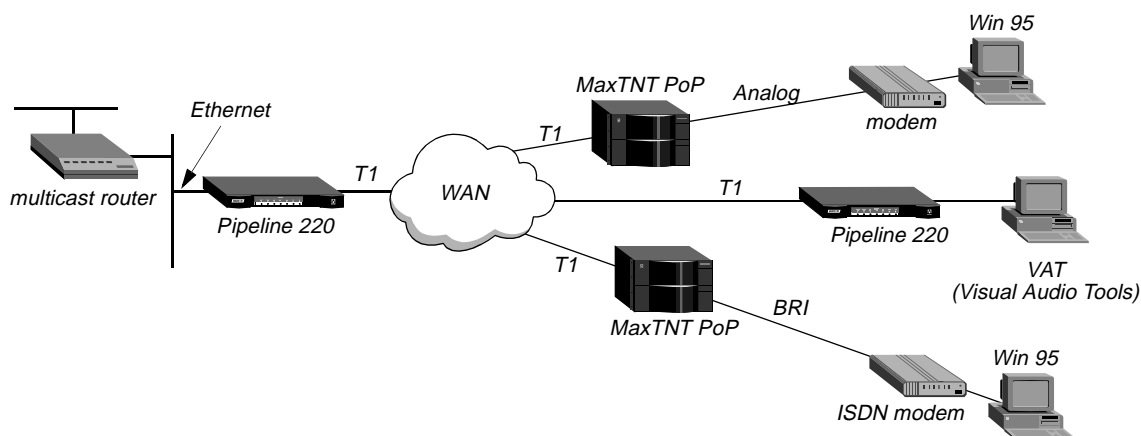


Figure 9-1. Pipeline 220 forwarding multicast traffic to multicast clients

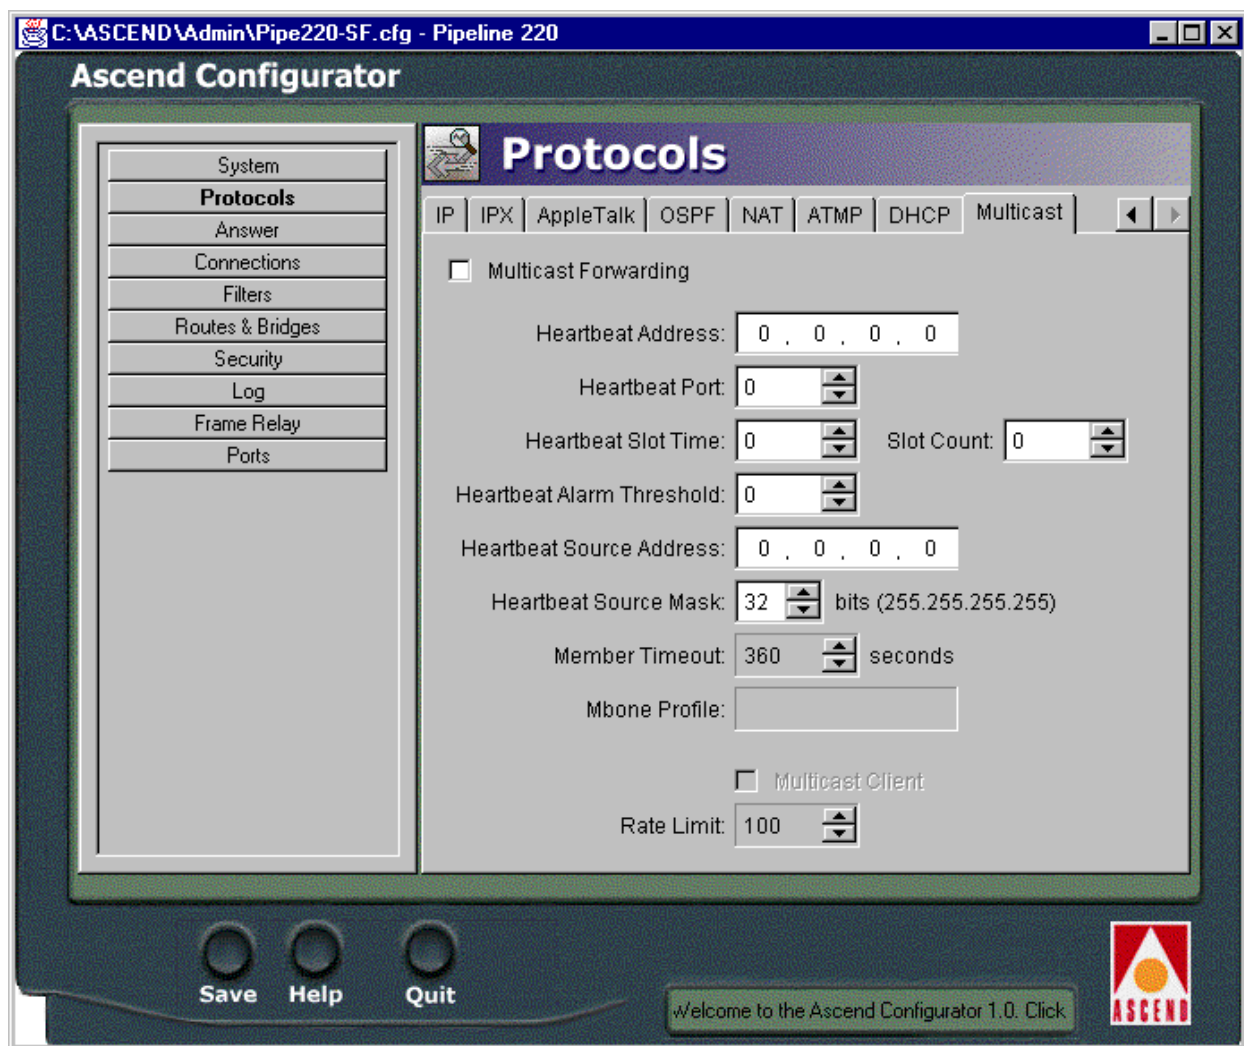
Note: Heartbeat monitoring is an optional feature. You can operate multicast forwarding without it if you prefer.

Configuring system-wide multicast parameters

The procedure configures system-wide multicast parameters. The values shown are examples that specify the MBONE interface as the Ethernet port, and uses the heartbeat group address of 224.1.1.1:

- 1 On the left side of the Configurator, click the Protocols button.

- 2 On the right side of the Configurator, click the Multicast tab.
Multicast parameters appear on the right side of the Configurator:



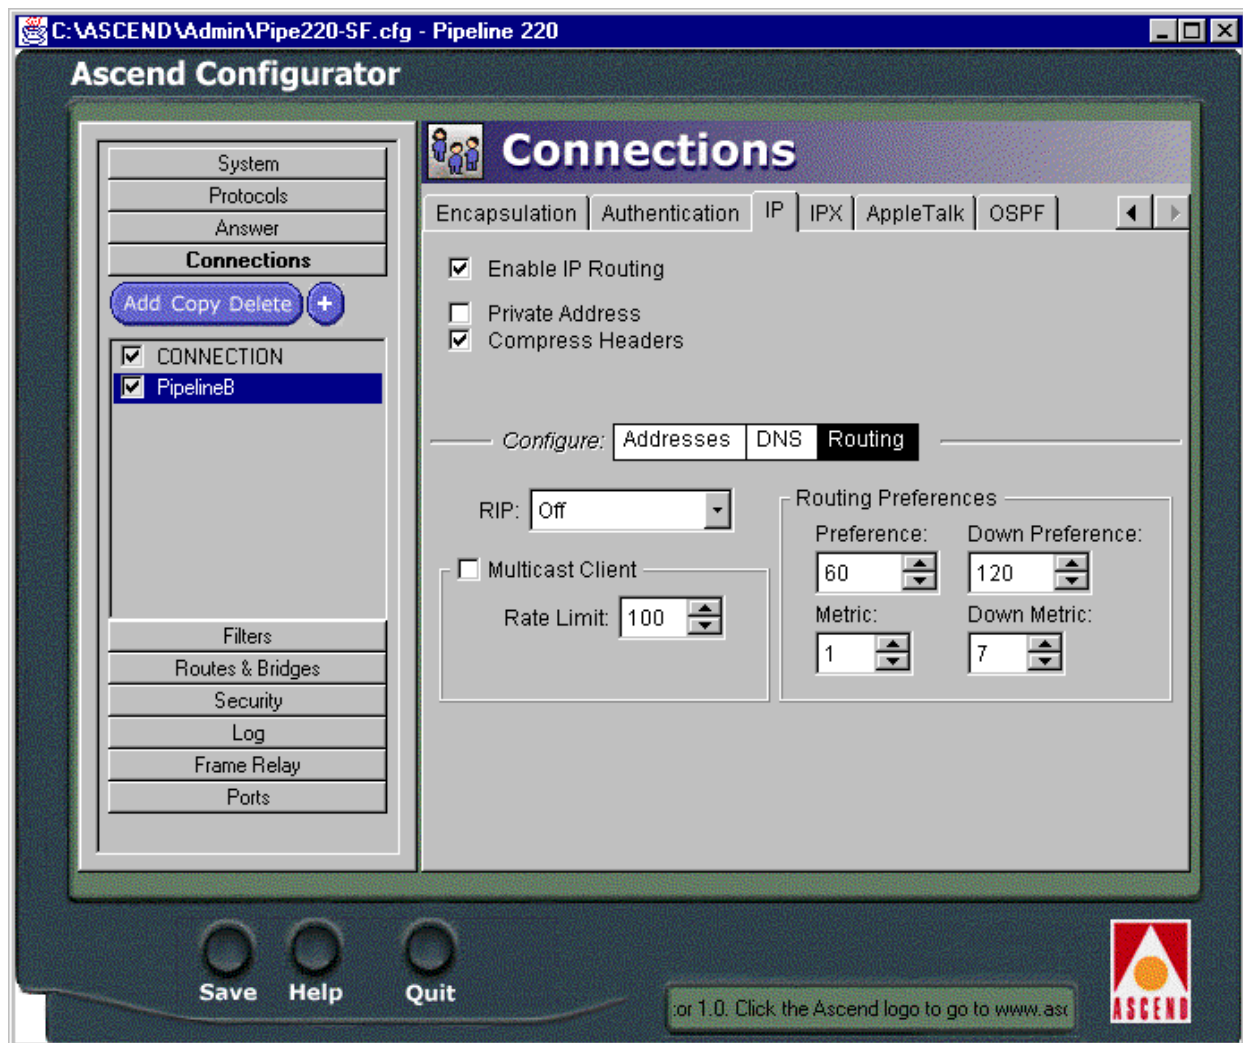
- 3 Select the Multicast Forwarding check box.
- 4 Specify Heartbeat Address.
To perform heartbeat monitoring, the Pipeline 220 looks for traffic destined for this address.
- 5 Specify Heartbeat Port.
To perform heartbeat monitoring, the Pipeline 220 looks for traffic destined for this port.
- 6 Set Heartbeat Slot Time to the number of seconds between Pipeline 220 polls for Multicast traffic. For example:
HeartBeat Slot Time=10
- 7 Set Slot Count to the quantity of polling cycles the Pipeline 220 waits before comparing the number of heartbeat packets received to the Alarm Threshold. For example:
HeartBeat Slot Count=10
- 8 Set Heartbeat Alarm Threshold.

The Pipeline 220 sends an SNMP alarm trap if the number of monitored packets falls below the Heartbeat Alarm Threshold.

Configuring multicasting on WAN interfaces

To enable multicasting on WAN interfaces:

- 1 On the left side of the Configurator, click the Connections button.
The Add/Copy/Delete window appears in the lower-left section of the Configurator.
- 2 From the Add/Copy/Delete window, highlight the connection.
The Configurator displays the parameters for the connection you select.
- 3 On the right side of the Configurator, click the IP tab.
The Configurator displays the parameters for the IP tab.
- 4 In the lower-right section of the Configurator, click the Routing button.
Multicast parameters appear in the lower-right section of the Configurator:



- 5 Select the Multicast Client check box.
- 6 Set Rate Limit to the rate at which the Pipeline 220 accepts multicast packets from clients on this interface. For example:

Rate Limit=5


To disable multicast forwarding on this interface, set Rate Limit to 100.

- 7 Enable the new settings, as described in “Saving the settings” on page 9-7.

Saving the settings

The Pipeline 220 settings you configure with the Ascend Configurator do not take effect until you upload them to the Pipeline 220. You use the Save command to upload the changes. Alternatively, you can save the settings to a text file for subsequent use, or to another device. To save the settings:

- 1 In the lower-left corner of the Configurator, click Save.
A dialog box appears, prompting you for a save method:



- 2 Select one of the following:
 - If you are uploading the configuration to the same device, select “Upload changes to <Pipeline 220 name>.”
 - If you are saving the configuration to a text file on your computer, select “Save a copy under a new filename” and enter the filename to which the configuration should be saved.
 - If you are uploading the configuration to a different device, select “Upload to the Ascend product at:” and enter the IP address of the device.
- 3 Click Save.

If you selected “Upload changes to <Pipeline 220 name>” in step 2, the Ascend Configurator uploads the configuration file to the Pipeline 220 via TFTP. If you selected one of the other options, the Configurator sends the file to the location you specified. If you specified a remote device, transmission is via TFTP.

Forwarding from an MBONE router on a WAN link

Figure 9-2 shows a multicast router on the WAN with local and multicast clients.

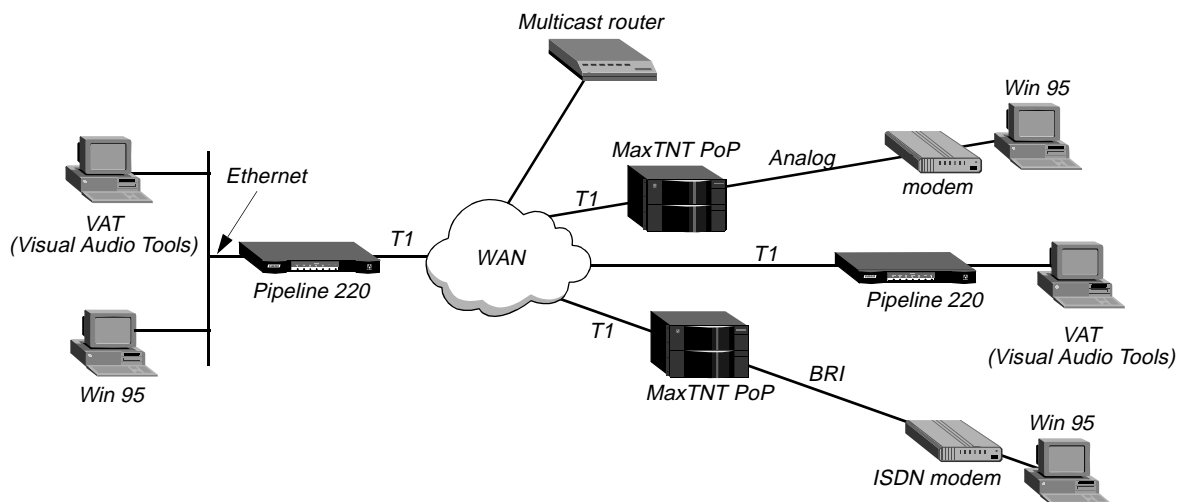


Figure 9-2. Pipeline 220 as a multicast forwarder on Ethernet and WAN interfaces

Note: This example does not use heartbeat monitoring. If you want to configure the Pipeline 220 for heartbeat monitoring, see the sample settings in “Configuring system-wide multicast parameters” on page 9-4.

This sample profile specifies the MBONE interface as a WAN link accessed through a Connection profile.

Configuring the Pipeline 220 to respond to multicast clients

To configure the Pipeline 220 to respond to multicast clients on the Ethernet:

- 1 On the left side of the Configurator, click the Protocols button.
- 2 On the right side of the Configurator, click the Multicast tab.
- 3 Select the Multicast Forwarding check box.
The Mbone Profile, Multicast Client, and Rate Limit parameters are now editable.
- 4 Set Mbone Profile to the Connection supporting MBONE.
- 5 Select the Multicast Client check box.
- 6 Set Rate Limit to a number lower than 100. For example,
Rate Limit=5
Setting it to 100 disables multicast forwarding on this interface.

Configuring the MBONE interface

To configure the MBONE interface:

- 1 On the left side of the Configurator, click the Connections button.
The Add/Copy/Delete window appears in the lower-left section of the Configurator.

- 2 In the Add/Copy/Delete window, highlight the connection to be used for the MBONE interface. You can also click Add to create a connection for the MBONE interface.
- 3 On the right side of the Configurator, click the IP tab.
- 4 In the lower-right section of the Configurator, click the Routing button.
- 5 Clear the Multicast Client check box.
- 6 Set Rate Limit to a value lower than 100. For example:
`Multicast Rate Limit=5`
Setting it to 100 disables multicast forwarding on this interface.

Configuring multicasting on WAN interfaces

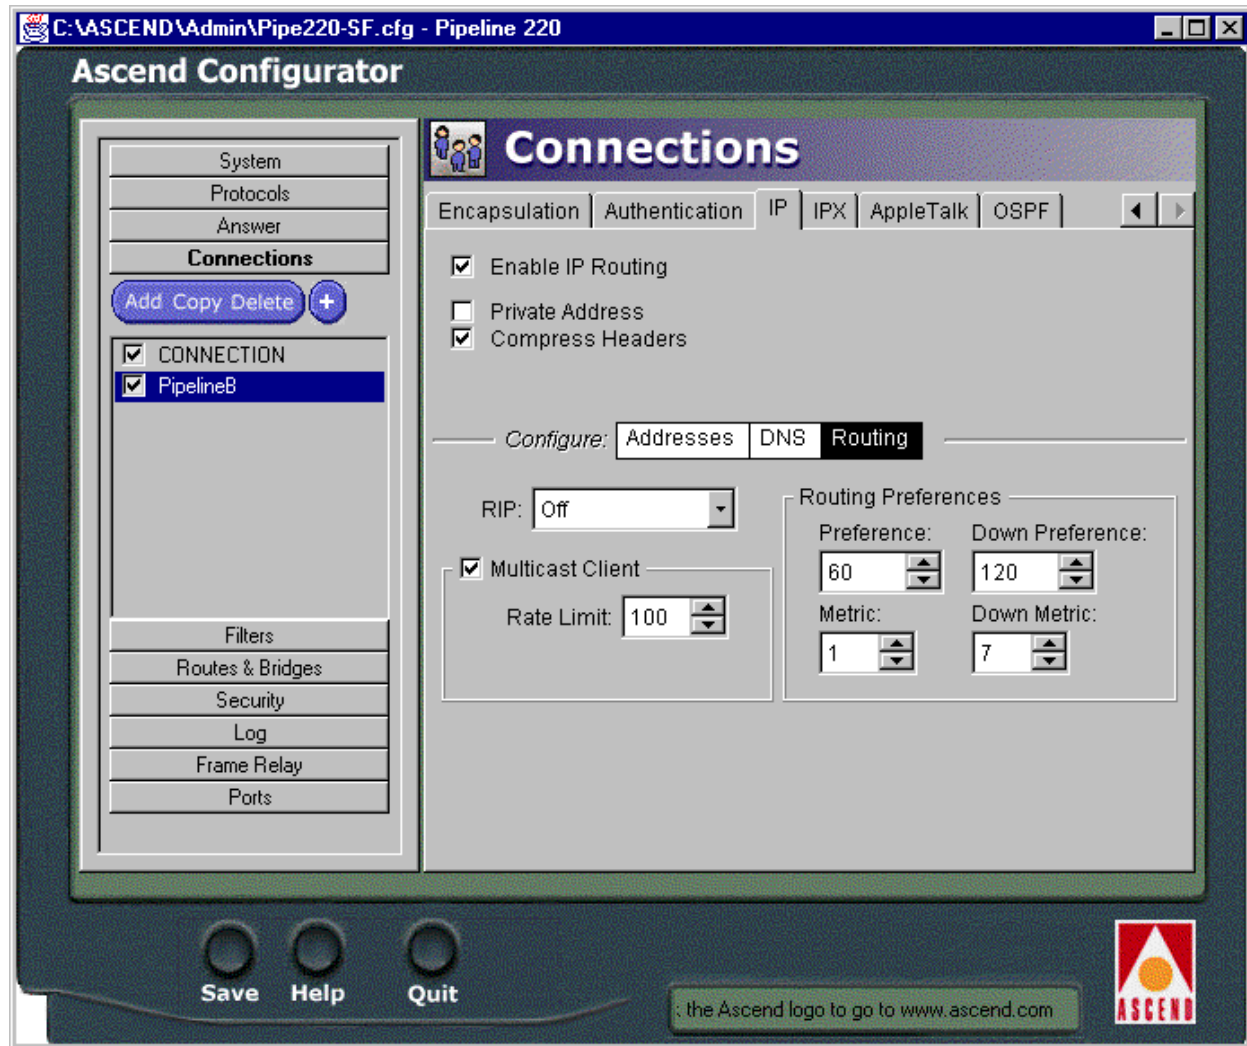
To enable multicasting on WAN interfaces, open the connection for a multicast client site:

- 1 On the left side of the Configurator, click the Connections button.
The Add/Copy/Delete window appears in the lower-left section of the Configurator.
- 2 In the Add/Copy/Delete window, highlight the connection to be used for the MBONE interface. You can also click Add to create a connection for the MBONE interface.
- 3 On the right side of the Configurator, click the IP tab.

Setting Up IP Multicast Forwarding

Forwarding from an MBONE router on a WAN link

- 4 In the lower-right section of the Configurator, click the Routing button.
Multicast parameters appear in the lower-center section of the Configurator:



- 5 Check the Multicast Client check box.
This step is the only difference from the configuration steps for the MBONE interface.
- 6 Set Rate Limit to a value lower than 100. For example:
Multicast Rate Limit=5
Setting it to 100 disables multicast forwarding on this interface.
- 7 Save the new settings, as described in “Saving the settings” on page 9-7.

Configuring IPX Routing

This chapter covers the following topics:

Introduction to Ascend IPX routing	10-1
Integrating the Pipeline 220 into the local IPX network	10-8
Working with the RIP and SAP tables	10-10
Example of an IPX routing connection.	10-19

Introduction to Ascend IPX routing

This section describes how the Pipeline 220 supports IPX routing between sites that run Novell NetWare version 3.11 or newer. The Pipeline 220 operates as an IPX router, with one interface on each of its two local Ethernet interfaces and the third across the WAN. Each IPX Connection profile defines an IPX WAN interface.

The most common use for IPX routing in the Pipeline 220 is to integrate multiple NetWare LANs to form an interconnected wide-area network

The Pipeline 220 supports IPX routing over PPP and Frame Relay connections. Support for both the IPXWAN and PPP IPXCP protocols makes the Pipeline 220 fully interoperable with non-Ascend products that conform to these protocols and associated RFCs.

Note: IPX transmission can use multiple frame types. The Pipeline 220, however, routes only one IPX frame type (which you configure), and it routes and spoofs IPX packets only if they are encapsulated in that frame. If bridging and IPX routing are enabled in the same Connection profile, the Pipeline 220 bridges any other IPX packet frame types. (For more information see Chapter 12, “Configuring Packet Bridging.”)

Unlike an IP routing configuration, where the Pipeline 220 uniquely identifies the calling device by its IP address, a Pipeline 220 IPX routing configuration does not include a built-in way to uniquely identify callers. For that reason, password authentication using PAP or CHAP is required unless IP routing is configured in the same Connection Profile.

IPX Service Advertising Protocol (SAP) tables

The Pipeline 220 follows standard IPX SAP behavior for routers. However, when it connects to another Ascend unit configured for IPX routing, the two units exchange their entire SAP tables. Each unit immediately adds all remote services to its SAP table.

NetWare servers broadcast SAP packets every 60 seconds to make sure that routers (such as the Pipeline 220) know about their services. Each router builds 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 its SAP-table entry for that server 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 220 consults its SAP table and replies with its own hardware address and the internal address of the requested server. This is analogous to proxy ARP in an IP environment.

Then the client transmits packets whose destination address is the internal address of the server. When the Pipeline 220 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 RIP (Routing Information Protocol) tables

The Pipeline 220 follows standard IPX RIP behavior for routers when connecting to non-Ascend units. However, when two Ascend units configured for IPX routing connect, they immediately exchange their entire RIP tables. In addition, the Pipeline 220 maintains those RIP entries as static until the unit is reset or power-cycled.

Note: In this chapter, RIP always refers to IPX RIP.

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

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 details.)

IPX routers broadcast RIP updates periodically and when a WAN connection is established. The Pipeline 220 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 220 recognizes network number -2 (0xFFFFFFF) as the IPX RIP default route. When it receives a packet for an unknown destination, it forwards the packet to the IPX router advertising the default route. For example, if the Pipeline 220 receives an IPX packet destined for network 77777777, and it does not have a RIP table entry for that destination, it forwards the packet towards network number FFFFFFFF, if available, instead of simply dropping the packet. If more than one IPX router is advertising the default route, the routing decision is based on Hop and Tick count.

Ascend extensions to standard IPX

NetWare uses dynamic routing and service location, so clients expect to be able to locate a server dynamically, regardless of where it is physically located. To help accommodate these expectations in a WAN environment, Ascend provides two IPX extensions: IPX Route profiles and IPX SAP filters.

(For information about the Handle IPX parameter and IPX bridging, see Chapter 12, “Configuring Packet Bridging.”)

IPX Route profiles

IPX Route profiles specify static IPX routes. When the Pipeline 220 clears its RIP and SAP tables because of a reset or power-cycle, it adds the static routes when it reinitializes. Each static route contains the information needed to reach one server.

If the Pipeline 220 connects to another Ascend unit, some sites choose not to configure a static route. Instead, after a power-cycle or reset, the initial connection to that site must be manually activated. After the initial connection, the Pipeline 220 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 help prevent timeouts when a client takes a long time to locate a server across a remote WAN link. (For more information, see “Configuring static IPX routes” on page 10-11, or see the *Configurator Online Help* for information about parameters in a profile.)

IPX SAP filters

Many sites do not want the Pipeline 220 SAP table to include long lists of all services available at a remote site. IPX SAP filters enable you to exclude services from, or explicitly include certain services in, the SAP table.

SAP filters can be applied to inbound or outbound SAP packets. Inbound filters control which services are added to the Pipeline 220 unit’s SAP table from advertisements on a network link. Outbound filters control which services the Pipeline 220 advertises on a particular network link. (For more information, see “Filtering SAP traffic” on page 10-15, or see the *Configurator Online Help* for information about parameters in a profile.)

WAN considerations for NetWare client software

NetWare clients on a wide-area network do not need special configuration in most cases. Following are some considerations regarding NetWare clients in an IPX routing environment, and Ascend’s recommendations.

Consideration	Recommendation
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 with the lowest connection costs. (See your NetWare documentation for more information.)
Local copy of LOGIN.EXE	Because of possible performance issues, executing programs remotely is not recommended. You should put LOGIN.EXE on each client’s local drive.

Consideration	Recommendation
Packet Burst (NetWare 3.11)	Packet Burst lets servers send a data stream across the WAN before a client sends an acknowledgment. The feature is enabled by default 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. (See your NetWare documentation for more information.)
Macintosh or UNIX clients	Both Macintosh and UNIX clients can use IPX to communicate with servers. But they also support native communications via AppleTalk or TCP/IP, respectively. If Macintosh clients must use AppleTalk software (rather than MacIPX) to access NetWare servers across the WAN, the WAN link must support bridging. Otherwise, AppleTalk packets will not make it across the connection. If UNIX clients access NetWare servers via TCP/IP (rather than UNIXWare), the Pipeline 220 must be configured as either a bridge or an IP router. Otherwise, TCP/IP packets will not make it across the connection.

IPX in the Answer profile

When the Pipeline 220 receives a request to bring up the WAN link, it checks the settings in its Answer profile. If the request does not include the information required by the Answer profile, the Pipeline 220 cannot successfully bring up the WAN link.

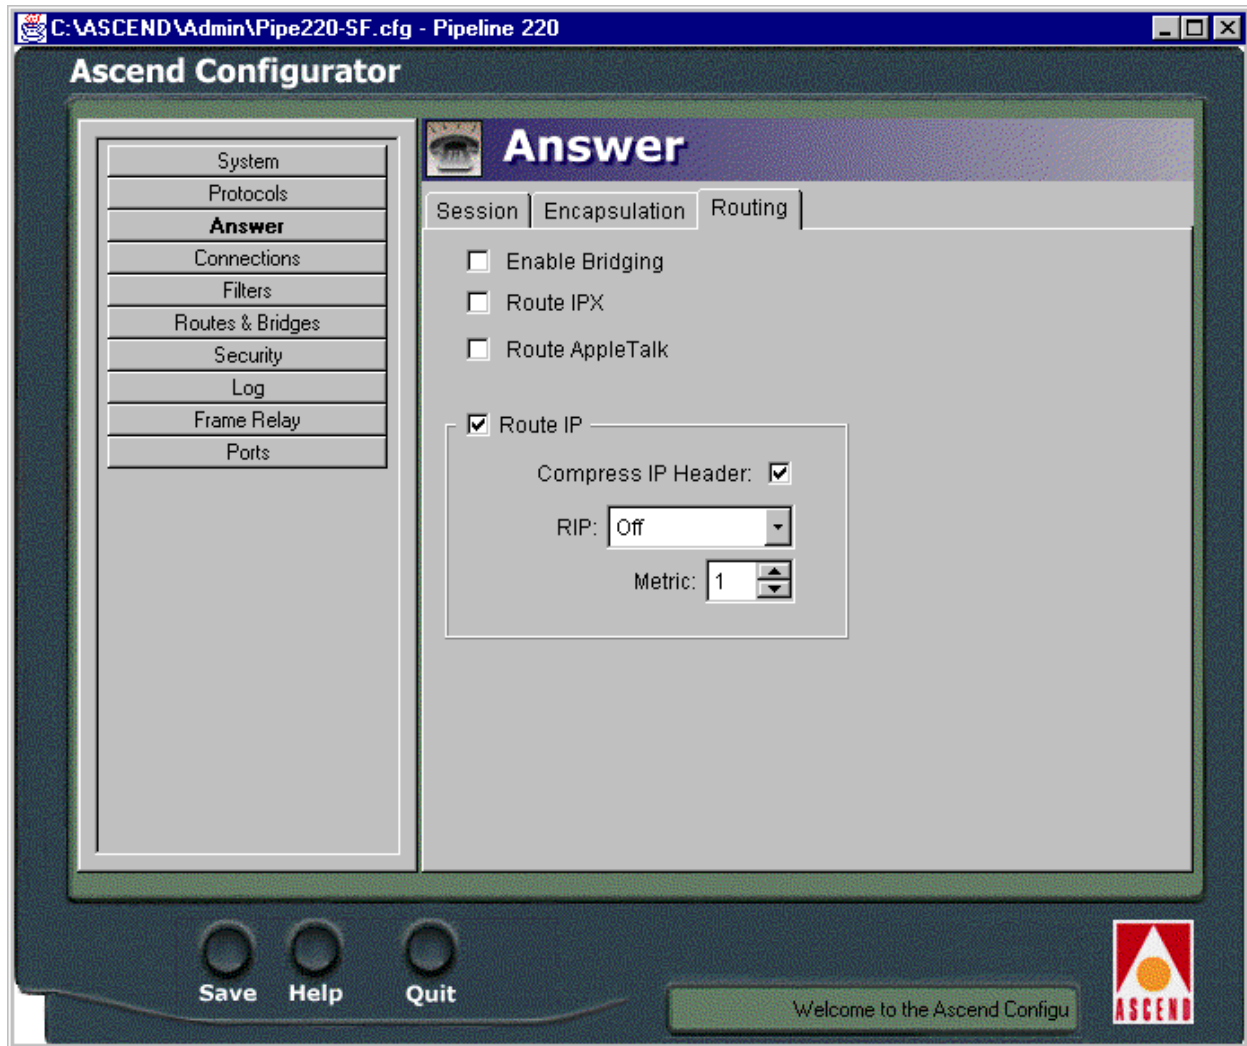
Note: Unlike an IP routing configuration, where the Pipeline 220 uniquely identifies the calling device by its IP address, an IPX routing configuration does not include a built-in way to uniquely identify callers. For that reason, password authentication using PAP or CHAP is required unless IP routing is configured in the same Connection profile.

Enabling IPX routing

To enable IPX routing:

- 1 On the left side of the Configurator, click the Answer button.

- 2 In the upper-right section of the Configurator, click the Routing tab.
System-wide bridging and routing parameters appear on the right side of the Configurator:



- 3 Select the Route IPX check box.

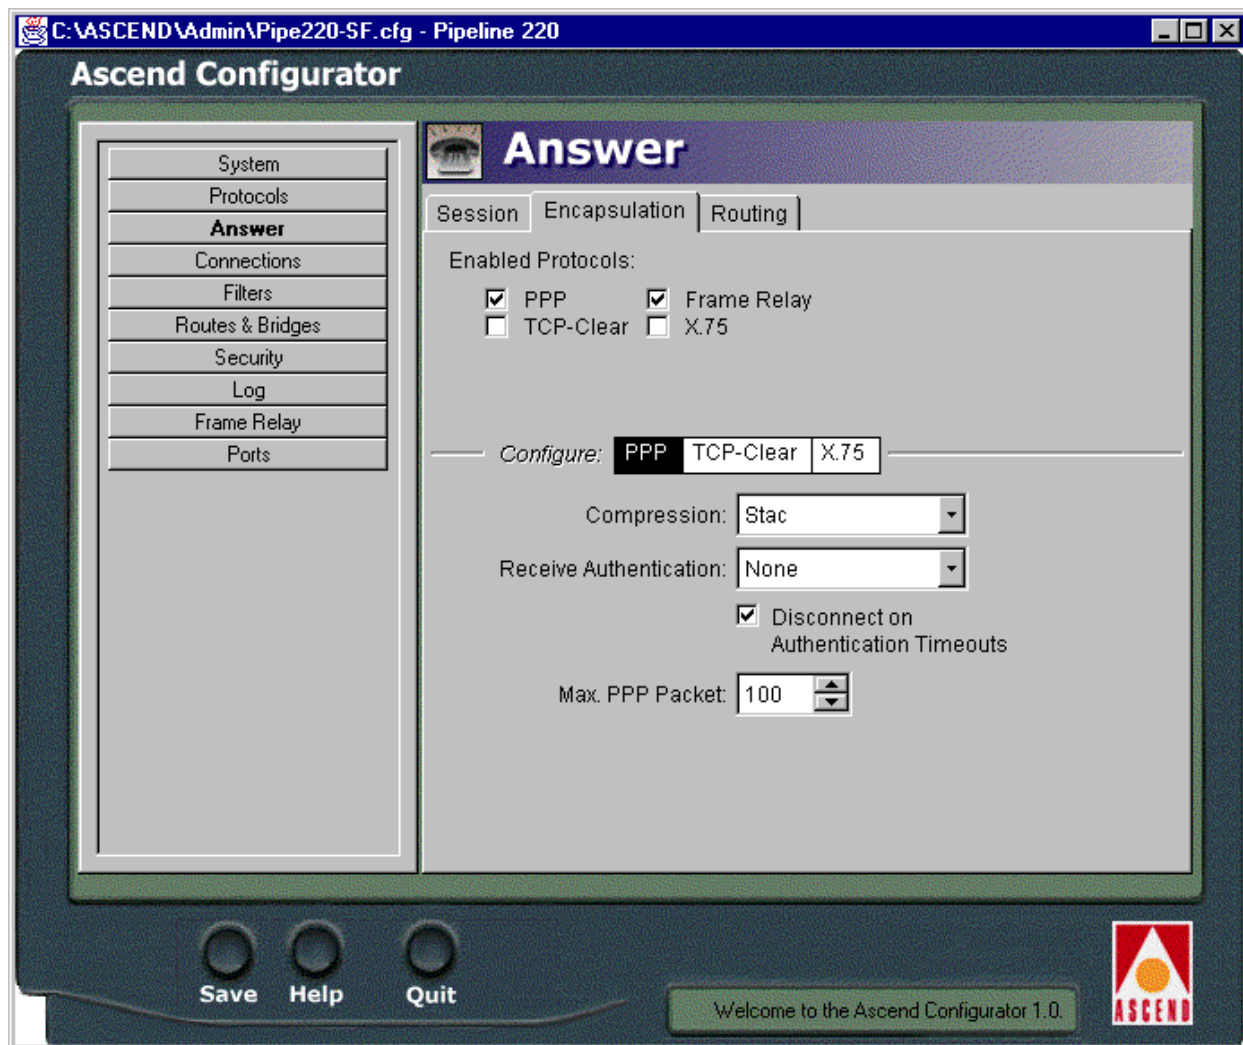
Enabling authentication

To enable password authentication for any remote request to bring up the WAN link:

- 1 In the upper-right section of the Configurator, click the Encapsulation tab.

- 2 Click the PPP button.

System-wide PPP parameters appear in the lower-right section of the Configurator:



- 3 Set Receive Authentication to Either.

Applying an IPX SAP Filter to the Answer profile

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

- 1 In the upper-right section of the Configurator, click the Sessions tab.
- 2 On the right section of the Configurator, click the Filters button.
- 3 Set the IPX SAP Filter parameter.

You apply an IPX SAP Filter profile by specifying the name you have given the IPX SAP filter. For example:

IPX SAP Filter=Office

For details, see “Filtering SAP traffic” on page 10-15.

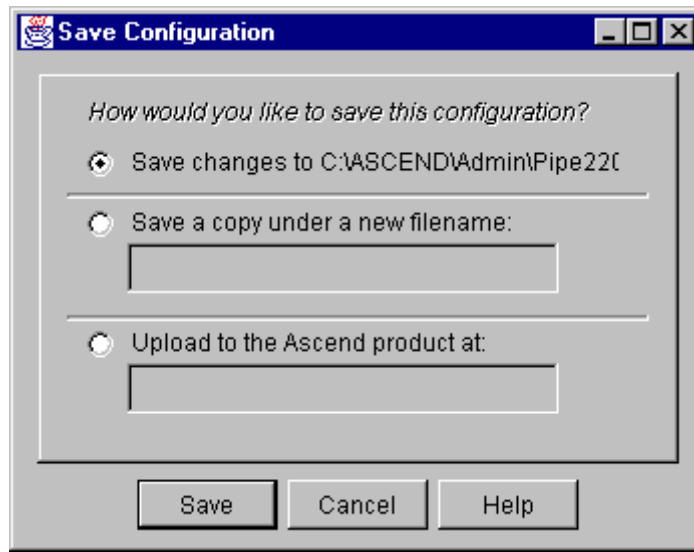
- 4 Save the new settings, as described in “Saving the settings” on page 10-7.

Saving the settings

The Pipeline 220 settings you configure with the Ascend Configurator do not take effect until you upload them to the Pipeline 220. You use the Save command to upload the changes. Alternatively, you can save the settings to a text file for subsequent use, or to another device.

To save the settings:

- 1 In the lower-left corner of the Configurator, click Save
A dialog box appears, prompting you for a save method:



- 2 Select one of the following:
 - If you are uploading the configuration to the same device, select “Upload changes to *<Pipeline 220 name>*.”
 - If you are saving the configuration to a text file on your computer, select “Save a copy under a new filename” and enter the filename to which the configuration should be saved.
 - If you are uploading the configuration to a different device, select “Upload to the Ascend product at:” and enter the IP address of the device.

- 3 Click Save.

If you selected “Upload changes to *<Pipeline 220 name>*” in step 2, the Ascend Configurator uploads the configuration file to the Pipeline 220 via TFTP. If you selected one of the other options, the Configurator sends the file to the location you specified. If you specified a remote device, transmission is via TFTP.

Integrating the Pipeline 220 into the local IPX network

To connect the Pipeline 220 to your local IPX network, you must perform the following tasks:

- Check any local NetWare server configurations to make sure that your Pipeline 220 configuration is consistent.
- On the Pipeline 220, enable IPX routing, and specify IPX frame type, and network number.

Checking local NetWare configurations

IPX packets are supported in more than one Ethernet frame type on an Ethernet segment. However, the Pipeline 220 can only route one IPX frame type (that you specify). It will, however, bridge any 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 the following:

```
internal network 1234
Bind ipx ipx-card net=CF0123FF
Load 3c509 name=ipx-card frame=ETHERNET_8023
```

In the output of the first line specifies the internal network number of the server. If you are not familiar with internal network numbers, see your NetWare documentation. Ascend units do not require internal network numbers.

The second line of output specifies the IPX network number in use on the Ethernet. The Pipeline 220 must use the same IPX network number for its Ethernet interface. You can specify the number explicitly in the Pipeline 220 Ethernet profile, or leave the Pipeline 220 number set to zero to enable it to learn the number from other routers.

The last 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 the Novell documentation.

Note: IPX network numbers on each network segment and internal network within a server on the *entire* WAN must have a unique network number. So you should know both the external and internal network numbers in use at all sites.

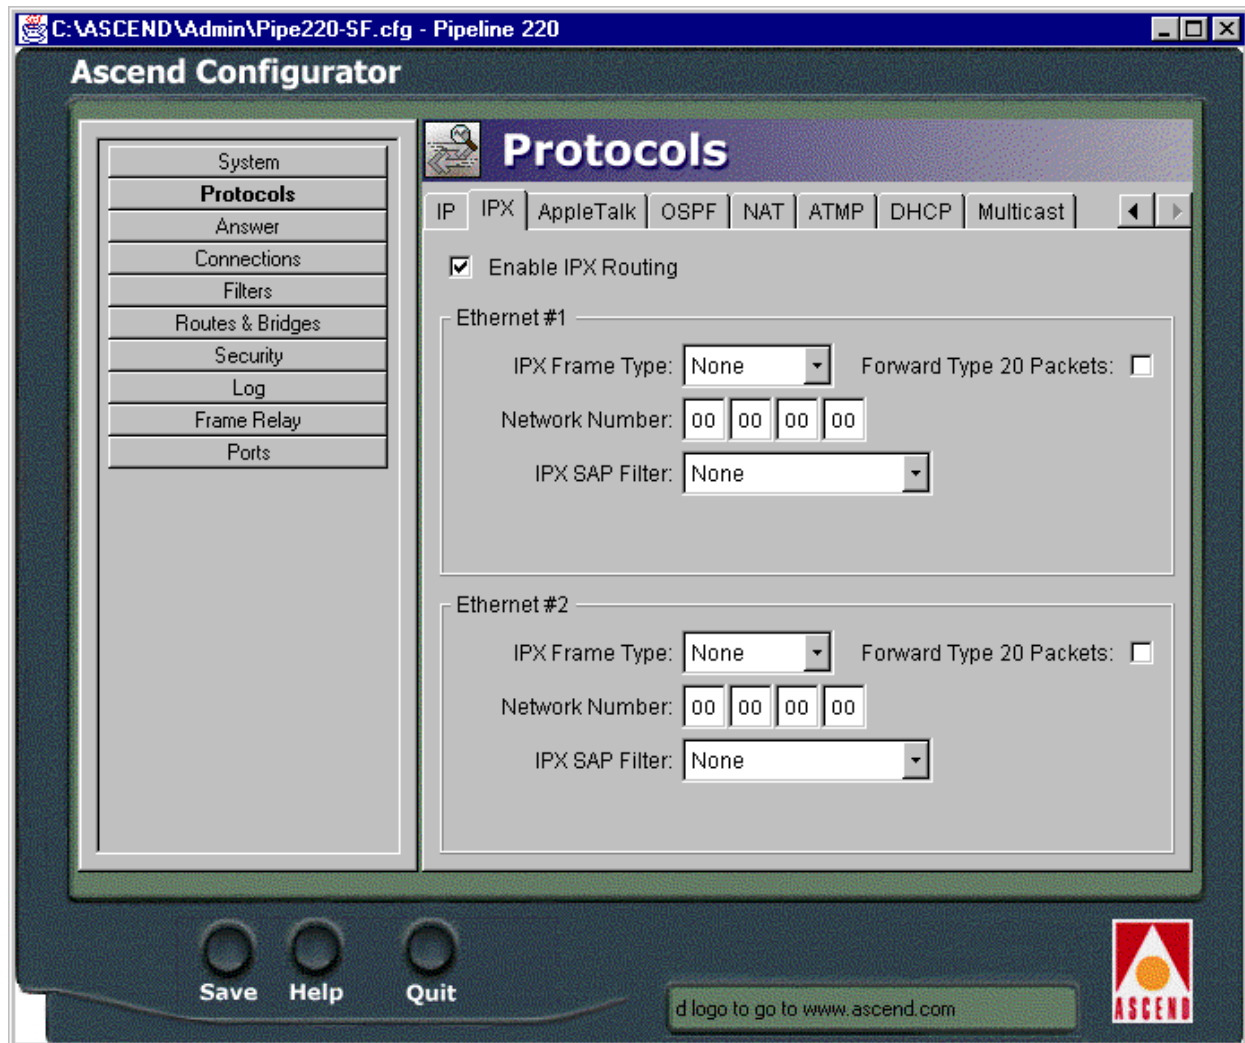
Configuring IPX on the Pipeline 220 Ethernet interface

By default, when you turn on IPX routing in the Pipeline 220 and close the Ethernet profile, the Pipeline 220 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 enable IPX routing in the Pipeline 220:

- 1 On the left side of the Configurator, click the Protocols button.

- 2 On the right side of the Configurator, click the IPX tab.
System-wide IPX parameters appear on the right side of the Configurator:



- 3 Select the Enable IPX check box.
- 4 Select the IPX frame type that the Pipeline 220 routes. 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.
- 5 Either configure the network number of the external network to which the Pipeline 220 is connected, or enable the Pipeline 220 to learn its IPX network number from other IPX local routers:
 - To configure the network number, make sure that the number is identical to the external network number of any other IPX router sharing the network cable with the Pipeline 220. If there are no other IPX routers sharing this network segment, be sure to choose a network number that is unique across the entire IPX internetwork.

- If IPX routers share the network cable with the Pipeline 220, you can set the network number to 00000000. This directs the Pipeline 220 to learn its address from other local routers.
- 6 Enable the new settings, as described in “Saving the settings” on page 10-7.

Working with the RIP and SAP tables

To manage 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 Connection profiles
- Configure a static route
- Configure SAP in Connection profiles
- Define and apply IPX SAP filters

You might also want to define standard call filters or data filters for additional control over WAN traffic and connections. (For details, see Chapter 13, “Defining Static Filters.”)

Viewing the RIP and SAP tables

You must use the Pipeline 220 terminal server to view the unit’s RIP and SAP tables. The terminal server is accessible only by using the VT100 interface. It is not available through the Configurator. (For additional information, see Chapter 16, “VT100 Interface System Administration.”)

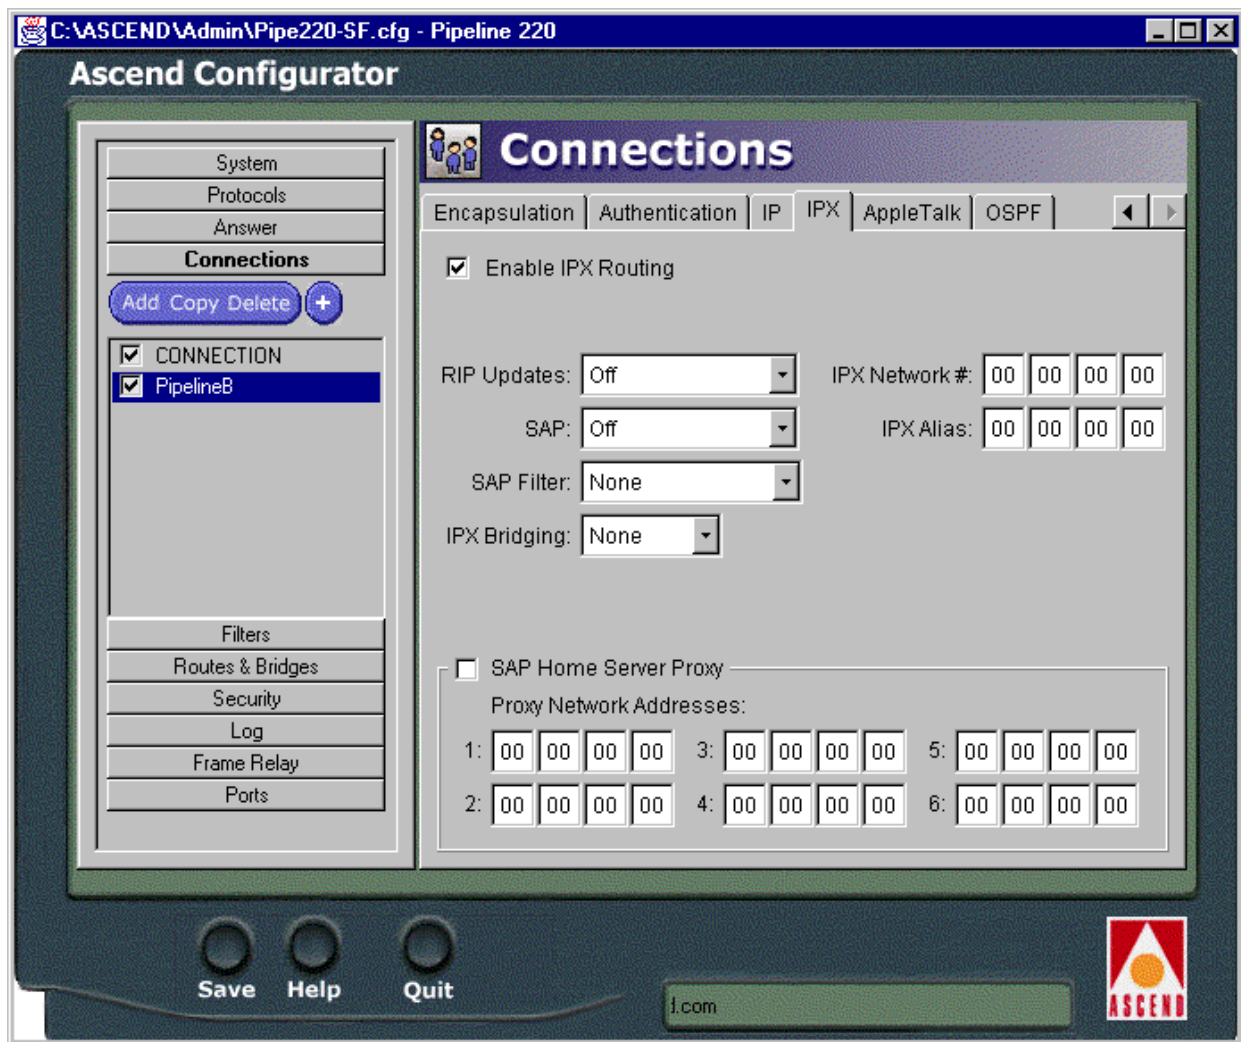
Restricting RIP in a Connection profile

By default, the IPX RIP parameter for a connection 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 220 only send or only receive RIP broadcasts on that connection. If the Pipeline 220 does not receive RIP broadcasts from a remote unit, you should configure a static route to at least one server on that network (described completely in the next section).

To restrict RIP exchange across a WAN connection:

- 1 On the left side of the Configurator, click the Connections button.
The Add/Copy/Delete window appears in the lower-left section of the Configurator.
- 2 In the Add/Copy/Delete window, highlight a connection that has IPX routing enabled.

- 3 On the right side of the Configurator, click the IPX tab.
IPX parameters for the connection appear on the right side of the Configurator:



- 4 Set the IPX RIP parameter to something other than its default value of Both.
For example:
`IPX RIP=Recv`
This setting means that the Pipeline 220 accepts RIP information from the remote IPX router but will not send its RIP information.
- 5 Enable the new settings, as described in “Saving the settings” on page 10-7.

Configuring static IPX routes

A static IPX route includes all of the information needed to reach one NetWare server on a remote network. When the Pipeline 220 receives an outbound packet for that server, it finds the referenced Connection profile and dials the connection. You configure the static route in an IPX Route profile.


You do not need to create IPX static routes to servers that are 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 the remote server.

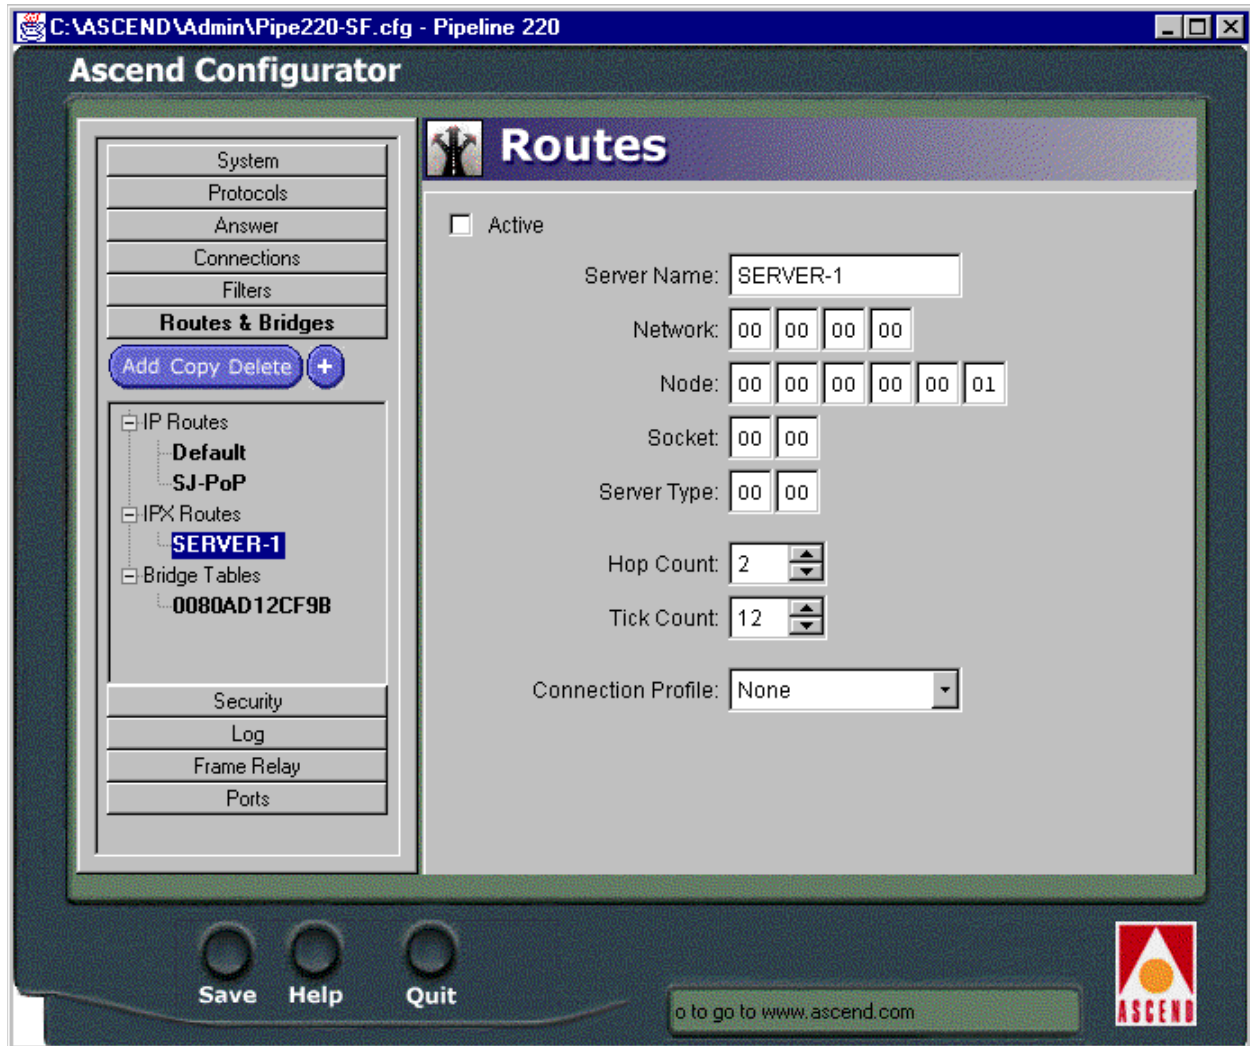
To define an IPX Route profile:

- 1 On the left side of the Configurator, click the Routes & Bridges button.
The Add/Copy/Delete window appears in the lower-left section of the Configurator.
- 2 In the Add/Copy/Delete window, highlight IPX Routes, then select the Add button.
The Configurator displays a dialog box prompting you for a name for the route:



- 3 Enter a name for the route and select OK.

The left side of the Configurator displays the new IPX static route. It has the name you entered, with default values for all other parameters:



- 4 Click the Active check box to specify that the route should be added to the RIP table.

- 5 Enter the remote server's internal network number. For example:

Network=ABC01FFF

- 6 Enter the remote server's node number. For example:

Node=00000000000001

The default 00000000000001 is typically the node number for NetWare file servers.

- 7 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 can use a different socket each time they load, and do not work with IPX Route profiles. To bring up a connection to a remote service that uses a dynamic socket

number, specify, on the remote network, a *master* server that uses a well-known socket number.

- 8** Specify the SAP Service Type.

For example:

```
Service Type=0004
```

NetWare file servers are SAP Service type 0004.

- 9** Specify the distance (in hops) to the server.

For example:

```
Hop count=2
```

Usually the default of 2 is appropriate.

- 10** Specify the distance to the server in IBM PC clock ticks (1/18 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.

- 11** Specify the name of the Connection profile that defines the WAN connection. For example:

```
Connection Name=TOREMOTE
```

- 12** Enable the new settings, as described in “Saving the settings” on page 10-7.

Restricting 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 220 broadcasts its SAP table to the remote network and listens for service updates from that network. Eventually, both networks have a full 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 220 will only send or only receive SAP broadcasts on that connection.

To restrict SAP across a WAN connection:

- 1** On the left side of the Configurator, click the Connections button.
The Add/Copy/Delete window appears in the lower-left section of the Configurator.
- 2** In the Add/Copy/Delete window, highlight a connection that has IPX routing enabled.
- 3** Set the IPX RIP parameter to something other than its default value of Both. For example:

```
IPX SAP=Recv
```

With this setting, the Pipeline 220 receives SAP table updates from the remote IPX router. If you do not want the Pipeline 220 to send or receive SAP broadcasts on this connection, set IPX SAP to None.
- 4** Enable the new settings, as described in “Saving the settings” on page 10-7.

Filtering SAP traffic

IPX SAP filters include or exclude specific NetWare services from the Pipeline 220 unit's SAP table. (Note that they do not help manage connectivity costs, unlike filters that prevent periodic RIP and SAP broadcasts from keeping a connection up unnecessarily.) 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). After you define an IPX SAP filter, you can apply it to an interface.

Defining an IPX SAP filter

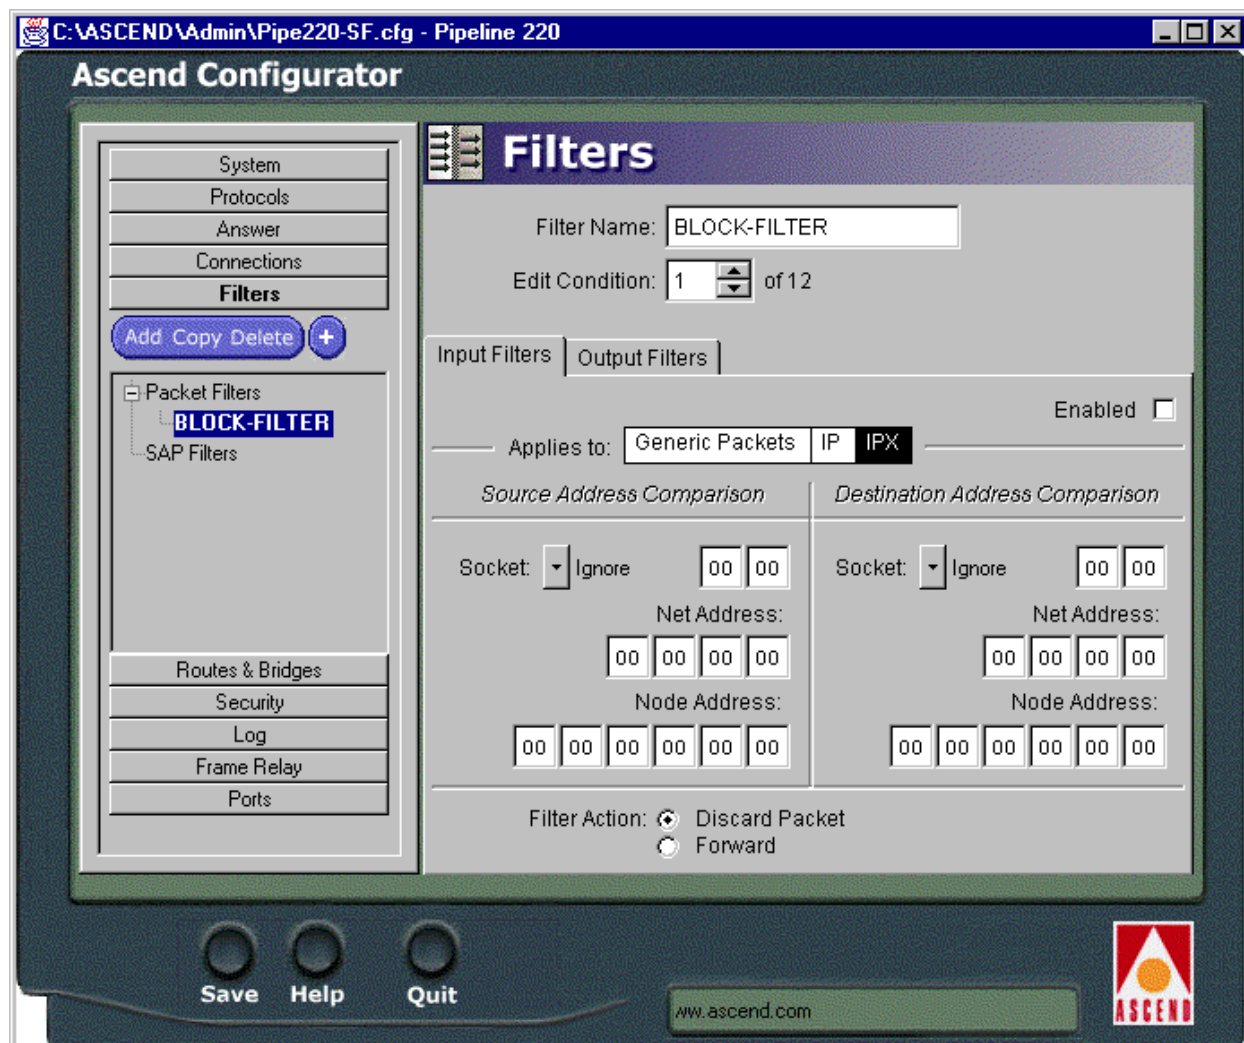
IPX SAP Input filters affect all packets the Pipeline 220 receives through the interface. Output filters affect all packets the Pipeline 220 sends through the interface.

To define an IPX SAP filter:

- 1** On the left side of the Configurator, click the Filters button.
The Add/Copy/Delete window appears in the lower-left section of the Configurator.
- 2** In the Add/Copy/Delete window, highlight SAP Filters, then select the Add button.
The Configurator displays a dialog box prompting you for a name for the SAP filter.

- 3 Enter a name for the SAP filter, and select OK.

The right side of the Configurator displays the new IPX SAP filter. It has the name you entered, with default values for all other parameters:



- 4 Set the parameters to define either an Input filter or an Output filter.

Defining an Input filter

Input filter conditions are applied to all SAP packets received by the Pipeline 220. They filter advertised services and exclude them from or include them in the Pipeline 220 SAP table.

You can specify up to 12 conditions for excluding services from or including services in the SAP table. Each of the 12 conditions can include or exclude either a type of service or a particular service. The Pipeline 220 applies these conditions in the order in which you list them: Input filter 1 followed by Input filter 2, and so forth. Proceed as follows:

- 1 Set Edit Condition to 1.
- 2 Click the Input Filters button.
- 3 Select the Enabled check box to activate Filter Condition 1.

- 4 Set Type to Exclude or Include, depending on whether you want to exclude or include a service.
- 5 Specify the service type (in hexadecimal format). For example:
`Server Type=4`
Note: File servers are service type 4.
- 6 Specify the NetWare server's name, as configured on the server.
- 7 If you want the Pipeline 220 to include or exclude other services, repeat step 1 through step 6 for the additional service. Set Edit Condition to 2. Repeat again as needed.
- 8 Enable the new settings, as described in "Saving the settings" on page 10-7.

Defining an output filter

Output filter conditions are applied to SAP response packets transmitted by the Pipeline 220. If the Pipeline 220 receives a SAP request packet, it applies Output filters before transmitting the SAP response, and excludes or includes services from the response packet as specified by the filter conditions.

You can specify up to 12 conditions for excluding services from or including services in the response packets. The Pipeline 220 applies the conditions in the order in which you list them: Output filter 1 followed by Output filter 2, and so forth.

- 1 Set Edit Condition to 1.
- 2 Click the Output Filters button.
- 3 Select the Enabled check box to activate Filter Condition 1.
- 4 Set Type to Exclude or Include, depending on whether you want to exclude or include a service.
- 5 Specify the service type (in hexadecimal format). For example:
`Server Type=4`
Note: File servers are service type 4.
- 6 Specify the NetWare server's name, as configured on the server.
- 7 If you want the Pipeline 220 to include or exclude other services, repeat step 1 through step 6 for the additional service. Set Edit Condition to 2. Repeat again as needed.
- 8 Enable the new settings, as described in "Saving the settings" on page 10-7.

Applying IPX SAP filters

You can apply an IPX SAP filter to the local Ethernet, to WAN interfaces, or both, to achieve the following effects:

- On Ethernet, a SAP filter includes or excludes specific servers or services from the table. If Directory Services are not supported, servers or services that are not in the Pipeline 220 table will be inaccessible to clients across the WAN.
- In the Answer profile, a SAP filter screens service advertisements from across the WAN if the remote device initiates the nailed-connection request.
- In a Connection profile, a SAP filter screens service advertisements from across the WAN if the P220 initiates the nailed-connection request.

(Although nailed connections do not function as switched connections, the initiation of the nailed connection is very similar to the initiation of a switched connection.)

Applying an IPX SAP filter to the Ethernet interface

To apply an IPX SAP filter to the Ethernet interface:

- 1** On the left side of the Configurator, click the Protocols button.
- 2** On the right side of the Configurator, click the IPX tab.
IPX options appear for both Ethernet interfaces.
- 3** For the Ethernet interface you are configuring, specify the name of the SAP filter you have defined.
- 4** Enable the new settings, as described in “Saving the settings” on page 10-7.

A filter applied to the Ethernet interface takes effect immediately.

Applying an IPX SAP filter to the Connection profile

To apply an IPX SAP filter to the Connection profile:

- 1** On the left side of the Configurator, click the Answer button.
- 2** On the right side of the Configurator, click the Sessions tab.
Sessions buttons appear in the lower-right section of the Configurator.
- 3** Click the Filters button.
- 4** Specify the name of the SAP filter you have defined.
- 5** Enable the new settings, as described in “Saving the settings” on page 10-7.

Example of an IPX routing connection

In this example, the Pipeline 220 is connected to an IPX network that supports both servers and clients and will connect with a remote site that also supports both servers and clients (as shown in Figure 10-1).

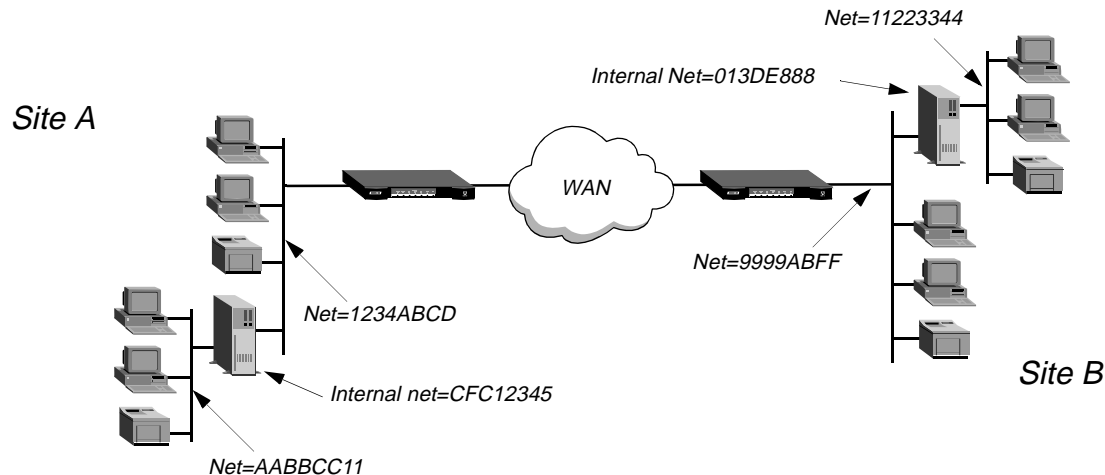


Figure 10-1. A connection with NetWare servers on both sides

Site A and site B are existing Novell LANs that support NetWare 3.12 and NetWare 4 servers, NetWare clients, and a Pipeline 220. 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
Bind ipx ipx-card2 net=AABBCC11
```

The NetWare server at site B is configured with the following information:

```
Name=SERVER-2
internal net 013DE888
Load 3c509 name=net-card frame=ETHERNET_8023
Bind ipx net-card net=9999ABFF
Bind ipx net-card2 net=11223344
```

To establish the connection shown in Figure 10-1, you would configure the Pipeline 220 at site A, enable IPX routing for its Ethernet interface, and configure a static route to the remote server. The same procedures would apply to site B.

Configuring the Pipeline 220 at site A

To configure the Pipeline 220 at site A, first specify the basic information about the connection, then create a Connection profile for connecting to site B:

- 1 Set a name for the Pipeline 220 in System button > Info tab > Name parameter.

```
Name=SITEAGW
```

Configuring IPX Routing

Example of an IPX routing connection

- 2 On the left side of the Configurator, click the Connections button.
The Add/Copy/Delete window appears in the lower-left side of the Configurator.
- 3 In the Add/Copy/Delete window, click Add.
The Configurator displays a dialog box prompting you for the name of the new connection.
- 4 Enter the name for the connection, then click OK. For example:
Station Name=TOSITEB
The connection profile appears on the right side of the Configurator, displaying the name and all other values at their factory default settings:



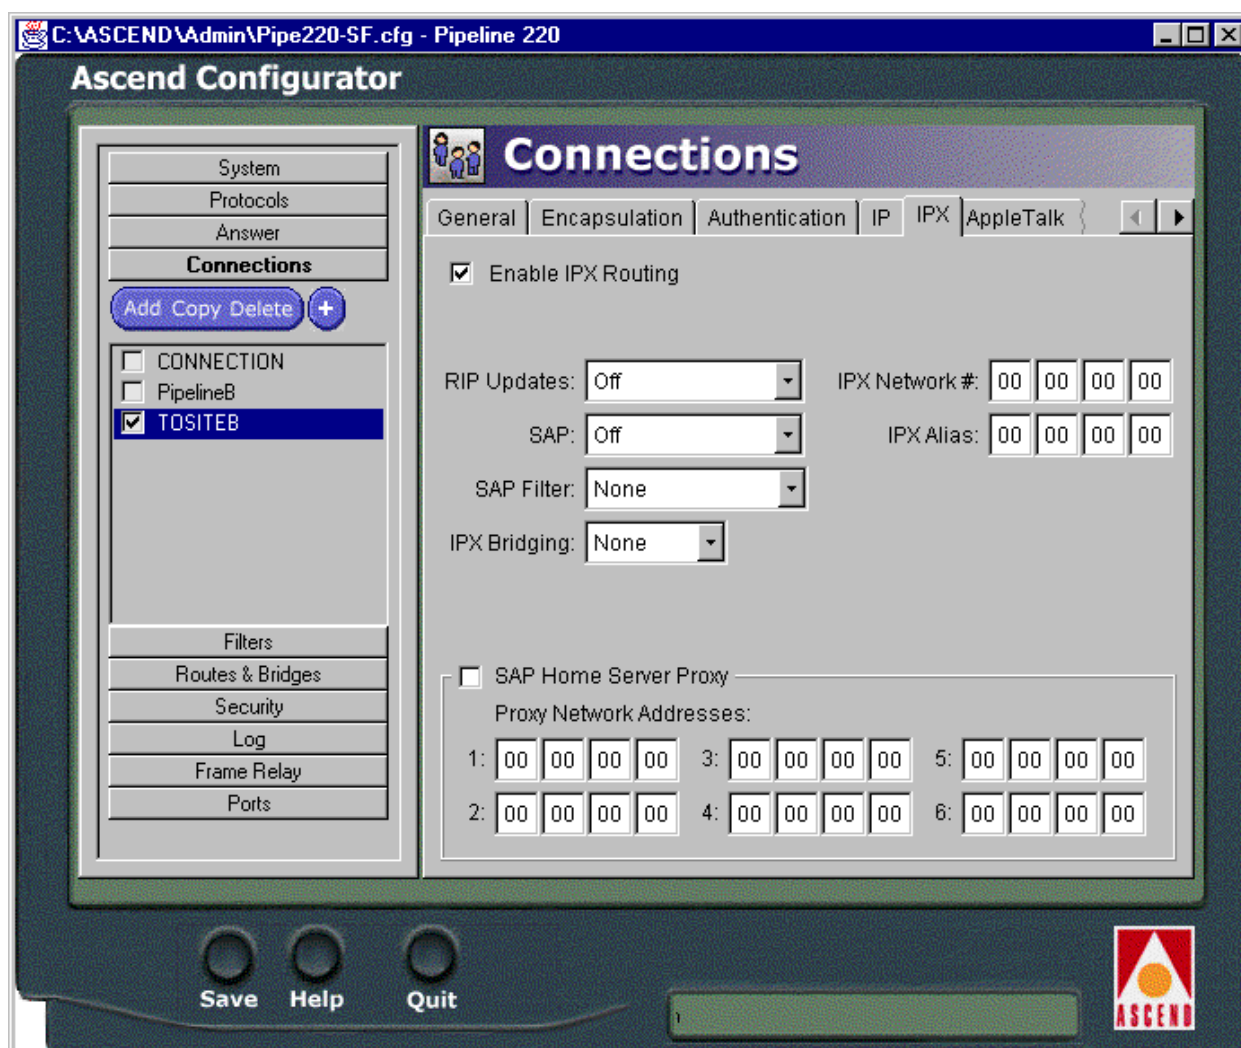
- 5 Enter the following parameter values:
 - Select the Connection button > General tab > Active check box.
 - Clear the Connection button > General tab > “Enable bridging” check box.
 - Clear the Connection button > IP tab > “Enable IP Routing” check box.
 - Select the Connection button > IPX tab > “Enable IPX Routing” check box.
 - Set Connection button > Authentication tab > Authentication=CHAP

- Set appropriate passwords in the Connection button > Authentication tab parameters.
 - Set Connection button > IPX tab > RIP Updates=None
 - Set Connection button > IPX tab > SAP=Both
- 6 Enable the new settings, as described in “Saving the settings” on page 10-7.

Enabling IPX routing for site A's Ethernet interface

To enable IPX routing for the Ethernet interface at site A:

- 1 On the left side of the Configurator, click the Protocols button.
 - 2 On the right side of the Configurator, click the IPX tab.
- System-wide IPX parameters appear for both Ethernet interfaces:



- 3 Select the Enable IPX Routing check box.
IPX routing is enabled for both Ethernet interfaces. (Clearing the check box disables it for both.)

Below the Enable IPX Routing check box are two boxes containing parameters for each Ethernet interface. For the Ethernet interface you are configuring:

4 Set IPX Frame Type=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.

5 Either set Network Number to the external network number of the segment to which the Pipeline 220 is connected, or enable the Pipeline 220 to learn its IPX network number from other IPX local routers:

- To configure the network number, make sure that the number is identical to the external network number of any other IPX router sharing the network cable with the Pipeline 220. If there are no other IPX routers sharing this network segment, be sure to choose a network number that is unique across the entire IPX internetwork.
- If IPX routers share the network cable with the Pipeline 220, you can set the network number to 00000000. This directs the Pipeline 220 to learn its address from other local routers.

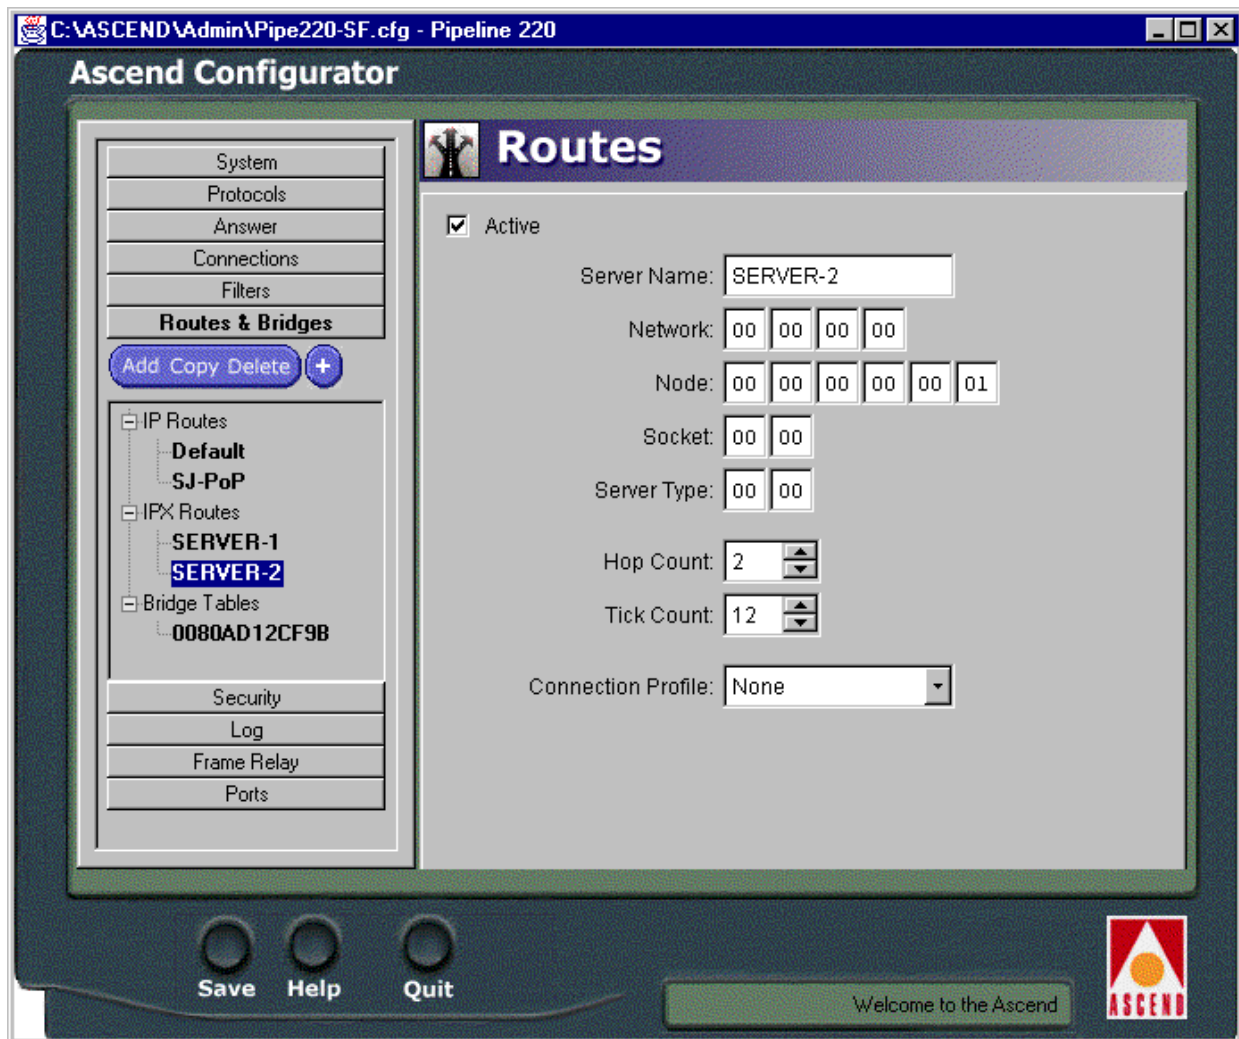
6 Enable the new settings, as described in “Saving the settings” on page 10-7.

Configuring a static route from site A to the remote server

Because IPX RIP is set to None in the Connection profile, configure a static route to the remote server:

- 1** On the left side of the Configurator, click the Routes & Bridges button.
The Add/Copy/Delete window appears in the lower-left section of the Configurator.
- 2** In the Add/Copy/Delete window, highlight IPX Routes, then select the Add button.
The Configurator displays a dialog box prompting you for a name for the filter.

- 3 For filter name, specify the name of the remote NetWare server, and select OK.
The left side of the Configurator displays the new IPX static route. It has the name you entered, with default values for all other parameters:



- 4 Specify that the route should be added to the RIP table:
Active=Yes
- 5 Enter the remote server's internal network number. For example:
Network=013DE888
- 6 Enter the remote server's node number:
Node=0000000000001
The default 0000000000001 is typically the node number for NetWare file servers.
- 7 Specify the remote server's socket number:
Socket=0451
Typically, Novell file servers use socket 0451.
- 8 Specify the SAP Service Type:
Service Type=0004

NetWare file servers are SAP Service type 0004.

- 9 Specify the distance (in hops) to the server:

`Hop count=2`

Usually the default of 2 is appropriate.

- 10 Specify the distance to the server in IBM PC clock ticks (1/18 second):

`Tick count=12`

Usually the default of 12 is appropriate, but you might need to increase this value for very distant servers.

- 11 Specify the name of the Connection profile that defines the WAN connection. For example:

`Connection Profile=TOSITEB`

Note: The Connection Name parameter in the IPX Route profile must match the name of the Connection profile you configured for the connection to that site.

- 12 Enable the new settings, as described in “Saving the settings” on page 10-7.

Configuring the Pipeline 220 at site B

To configure the Pipeline 220 at site B, first specify the basic information about the connection:

- 1 Set a name for the Pipeline 220 in System button > Info tab > Name parameter.

`Name=SITEBGW`

- 2 On the left side of the Configurator, click the Connections button.

The Add/Copy/Delete window appears in the lower-left section of the Configurator.

- 3 Create a Connection profile for connecting to site A.

- 4 In the Add/Copy/Delete window, click Add.

The Configurator displays a dialog box prompting you for a name for the new connection.

- 5 Enter name for the connection, then click OK.

`Station Name=TOSITEA`

The connection profile appears on the right side of the Configurator, displaying the name and all other values at their factory default settings.

- 6 Set the following parameters:

- Select the Connection button > General tab > Active check box.
- Clear the Connection button > General tab > “Enable bridging” check box.
- Clear the Connection button > IP tab > “Enable IP Routing” check box.
- Select the Connection button > IPX tab > “Enable IPX Routing” check box.
- Set Connection button > Authentication tab > Authentication=CHAP.
- Set appropriate passwords in the Connection button > Authentication tab parameters.
- Set Connection button > IPX tab > RIP Updates=None.
- Set Connection button > IPX tab > SAP=Both.

- 7 Enable the new settings, as described in “Saving the settings” on page 10-7.

Enabling IPX routing for site B's Ethernet interface

Enable IPX routing for the Ethernet interface at site B:

- 1 On the left side of the Configurator, click the Protocols button.
- 2 On the right side of the Configurator, click the IPX tab.
IPX-related parameters for both Ethernet interfaces appear.
- 3 Select the Enable IPX Routing check box.
IPX routing is enabled for both Ethernet interfaces. (Clearing the check box disables it for both interfaces.)

Below the Enable IPX Routing check box are two boxes containing parameters for each Ethernet interface. For the Ethernet interface you are configuring:

- 4 Set IPX Frame Type=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.
- 5 Either set Network Number to the external network number of the segment to which the Pipeline 220 is connected, or enable the Pipeline 220 to learn its IPX network number from other IPX local routers:
 - To configure the network number, make sure that the number is identical to the external network number of any other IPX router sharing the network cable with the Pipeline 220. If there are no other IPX routers sharing this network segment, be sure to choose a network number that is unique across the entire IPX internetwork.
 - If IPX routers share the network cable with the Pipeline 220, you can set the network number to 00000000. This directs the Pipeline 220 to learn its address from other local routers.
- 6 Enable the new settings, as described in “Saving the settings” on page 10-7.

Configuring a static route at site B

Because IPX RIP is set to None in the Connection profile, configure a static route to the remote server:

- 1 On the left side of the Configurator, click the Routes & Bridges button.
The Add/Copy/Delete window appears in the lower-left section of the Configurator.
- 2 In the Add/Copy/Delete window, highlight IPX Routes, then select the Add button.
The Configurator displays a dialog box prompting you for a name for the route.
- 3 Enter a name for the route, and select OK.
The left side of the Configurator displays the new IPX static route. It has the name you entered, with default values for all other parameters.
- 4 Specify the name of the remote NetWare server. For example:
`Server Name=SERVER-1`
- 5 Specify that the route should be added to the RIP table.
`Active=Yes`
- 6 Enter the remote server's internal network number. For example:
`Network=CFC12345`

- 7 Enter the remote server's node number:
`Node=0000000000001`
The default 0000000000001 is typically the node number for NetWare file servers.
- 8 Specify the remote server's socket number:
`Socket=0451`
Typically, Novell file servers use socket 0451.
- 9 Specify the SAP Service Type:
`Service Type=0004`
NetWare file servers are SAP Service type 0004.
- 10 Specify the distance (in hops) to the server:
`Hop count=2`
Usually the default of 2 is appropriate.
- 11 Specify the distance to the server in IBM PC clock ticks (1/18 second):
`Tick count=12`
Usually the default of 12 is appropriate, but you might need to increase this value for very distant servers.
- 12 Specify the name of the Connection profile that defines the WAN connection. For example:
`Connection Profile=TOSITEA`

Note: The Connection Name parameter in the IPX Route profile must match the name of the Connection Profile you configured to that site.
- 13 Enable the new settings, as described in "Saving the settings" on page 10-7.

Configuring AppleTalk Routing

This chapter covers the following topics:

Introduction to AppleTalk routing	11-1
How AppleTalk works	11-4
Configuring AppleTalk routing	11-5

Introduction to AppleTalk routing

The Pipeline 220 functions as an AppleTalk internet router, providing routing functions for AppleTalk nodes (Macintosh workstations or Apple printers) that are connected to the Pipeline 220 over Ethernet or a WAN. The following AppleTalk protocols are supported:

- Datagram Delivery Protocol (DDP)
- Routing Table Maintenance Protocol (RTMP)
- AppleTalk Echo Protocol (AEP)
- Zone Information Protocol (ZIP)
- Name Binding Protocol (NBP)
- AppleTalk Control Protocol (ATCP— for router-to-router applications)

When to use AppleTalk routing

With AppleTalk routing, connect two or more networks that have AppleTalk nodes, such as Mac OS computers or Apple printers. The primary benefits of routing AppleTalk traffic (as opposed to bridging this traffic) are:

- Reducing broadcast and multicast traffic over the WAN
- Providing startup information to local AppleTalk devices

Reducing broadcast and multicast traffic

Because AppleTalk uses multicast and broadcast addresses extensively, routing AppleTalk can greatly improve the efficiency of a LAN or WAN. By using AppleTalk zones to segment traffic, you can significantly reduce the amount of broadcast and multicast traffic on a LAN or WAN. When you set up a router for the first time, you identify the cable range (network-number range) for the subnetwork segment and one or more zones.

For example, when a user on a network without a router selects a device in the Chooser, the MAC OS computer sends out a Name Binding Protocol (NBP) Lookup as a broadcast packet. Since a bridge forwards all broadcast traffic, all devices on the network receive the Lookup

packet. A router can significantly reduce AppleTalk traffic over the WAN because it does not forward broadcast traffic from one subnetwork to another, but stops it at the subnetwork port of the router.

Zone multicasting is intended to prevent any node not in the destination zone for the lookup from receiving the lookup packet. Any AppleTalk node responds only to NBP lookups for that node's zone name. In the example above, a router would convert the Broadcast Request packet generated by the Lookup request to a Forward Request packet for each network that contains nodes in the target zone specified by the Lookup request.

A bridge can filter directed traffic between two specific nodes but cannot filter broadcast or multicast traffic, since there isn't a specific port that can be assigned to a multicast or broadcast address. This means that although filters used with bridging can reduce the number of AppleTalk packets sent to remote network segments, bridging does not reduce the number of broadcast and multicast packets over these networks.

Providing dynamic startup information to local devices

In addition to routing services, the Ascend AppleTalk router provides startup information to AppleTalk stations. As with other routed protocols, AppleTalk station, or *node*, addresses are comprised of a unique network number/node combination. AppleTalk addresses are dynamically assigned when a node starts up. In addition, the router provides an AppleTalk node with the network cable range to which it is attached, and supplies zone name information.

Understanding AppleTalk zones and network ranges

AppleTalk zones and network ranges are configured in AppleTalk routers. Network numbers are assigned to network segments, and must be unique within the internetwork. A network range is a range of network numbers set into the port descriptor of the router port and then transmitted through RTMP to the other nodes of the network. Each of the numbers within a network range can represent up to 253 devices.

AppleTalk zones

A zone is a multicast address containing an arbitrary subset of the AppleTalk nodes in an internet. Each node belongs to only one zone, but a particular extended network can contain nodes belonging to any number of zones. Zones provide departmental or other groupings of network entities that a user can easily understand.

In the Ascend AppleTalk router, zone names are case-insensitive. However, since some routers regard zone names as case-sensitive, it is advisable to be consistent in spelling zone names when you configure multiple connections or routers.

Extended and non-extended AppleTalk networks

AppleTalk subnetworks are either non-extended or extended. Non-extended networks theoretically allow up to 254 nodes. A non-extended network has one network number (not a range) and one zone. Examples of non-extended networks are LocalTalk and ARA dial-up networks.

An extended network is a group of non-extended networks on the same physical data link, and contains a range of network numbers. Each network in the range supports up to 253 devices. EtherTalk and TokenTalk are examples of extended networks.

At least one router on a network, called the seed router, must have the network number range set into its port description. Other routers on the network can have a network range of 0 (zero), which specifies that they acquire the network-number range from RTMP packets sent by the seed router. AppleTalk routers on a network must not have conflicting network-number ranges for that network. A 0 value does not cause a conflict, but otherwise, all seed routers on the same network must have the same value for the start and end of the network-number range.

Figure 11-1 shows a network with three routers and three zones configured. Each zone has a range of network numbers.

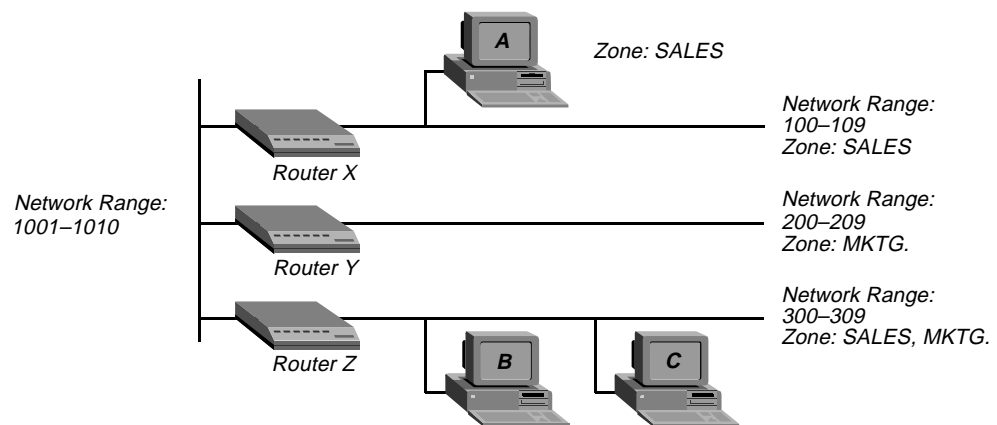


Figure 11-1. AppleTalk LAN

Router X, Router Y, and Router Z connect to the backbone network (Range 1001-1010). Each router has an additional connection to a local network segment. For example, Router X has a connection to the network range 100-109. User A's computer also connects to the 100-109 range.

Because Router X is configured with only one zone, any AppleTalk device joining the segment belongs to the SALES zone. But User B's computer can belong to either the SALES zone or the MKTG. zone. Some AppleTalk devices allow you to select the zone to which they belong. If there is no way to manually assign the zone, the AppleTalk device is put into the *default* zone, which is defined on the AppleTalk router.

Figure 11-1 shows two important concepts about network numbers and zones. When a network range is defined, all values within that range are unusable for any other segment. The segment to which user C's computer connects uses network range 300-309. No other network segment in this AppleTalk network can use network numbers 300, 301, 302, etc., in their ranges. As an example, network number 310 is available to a new network segment

Zones can be shared among network segments. In Figure 11-1, network 100-109 supports zone SALES. So does network 300-309.

How AppleTalk works

The following is a brief description of how the workstation user sees a typical AppleTalk connection and describes in a general way what is happening as the user makes the choices that lead to a connection. This example supposes a connection between a workstation on a Pipeline 220 connected to another Pipeline 220 over Ethernet on a synchronous PPP connection, as shown in Figure 2.




Figure 11-2. Routed connection

- 1 An AppleTalk workstation user opens the Chooser for the first time since it has been attached to the router and configured.
The zones that appear are in the local Ethernet zone (in this case the WAN zone is the same as the local Ethernet zone), configured in the Connection profile for the Pipeline 220. This information is stored in the Pipeline 220.
- 2 The workstation sends a ZIP Query to obtain an updated zone list from the Pipeline 220, and the Pipeline 220 returns the updated zone list. This list might contain different zones than did the initial list.
- 3 The user selects a zone and a specific device in the Chooser.
- 4 The workstation sends an NBP Broadcast Request to the Pipeline 220, which checks its Zone Information Table to determine which subnetwork that printer is located in, and sends the request to the other Pipeline 220 via the port configured in the Connection Profile.
- 5 The remote Pipeline 220 determines the port to which the subnetwork is attached and performs the lookup in the appropriate multicast address (multicast addresses are assigned to zones).
- 6 All devices in the appropriate zone on the subnetwork hear and pick up the NBP Lookup.
- 7 The selected printer obtains the sender's address from the Lookup packet (in this case the routers are *forwarders*; the workstation is the *sender*) and sends the reply through the routers to the workstation.
- 8 The user sends the print job to the printer.
- 9 When the print job is complete and no data packets are passing through the connection, the Pipeline 220s continue to pass routing information.

Configuring AppleTalk routing

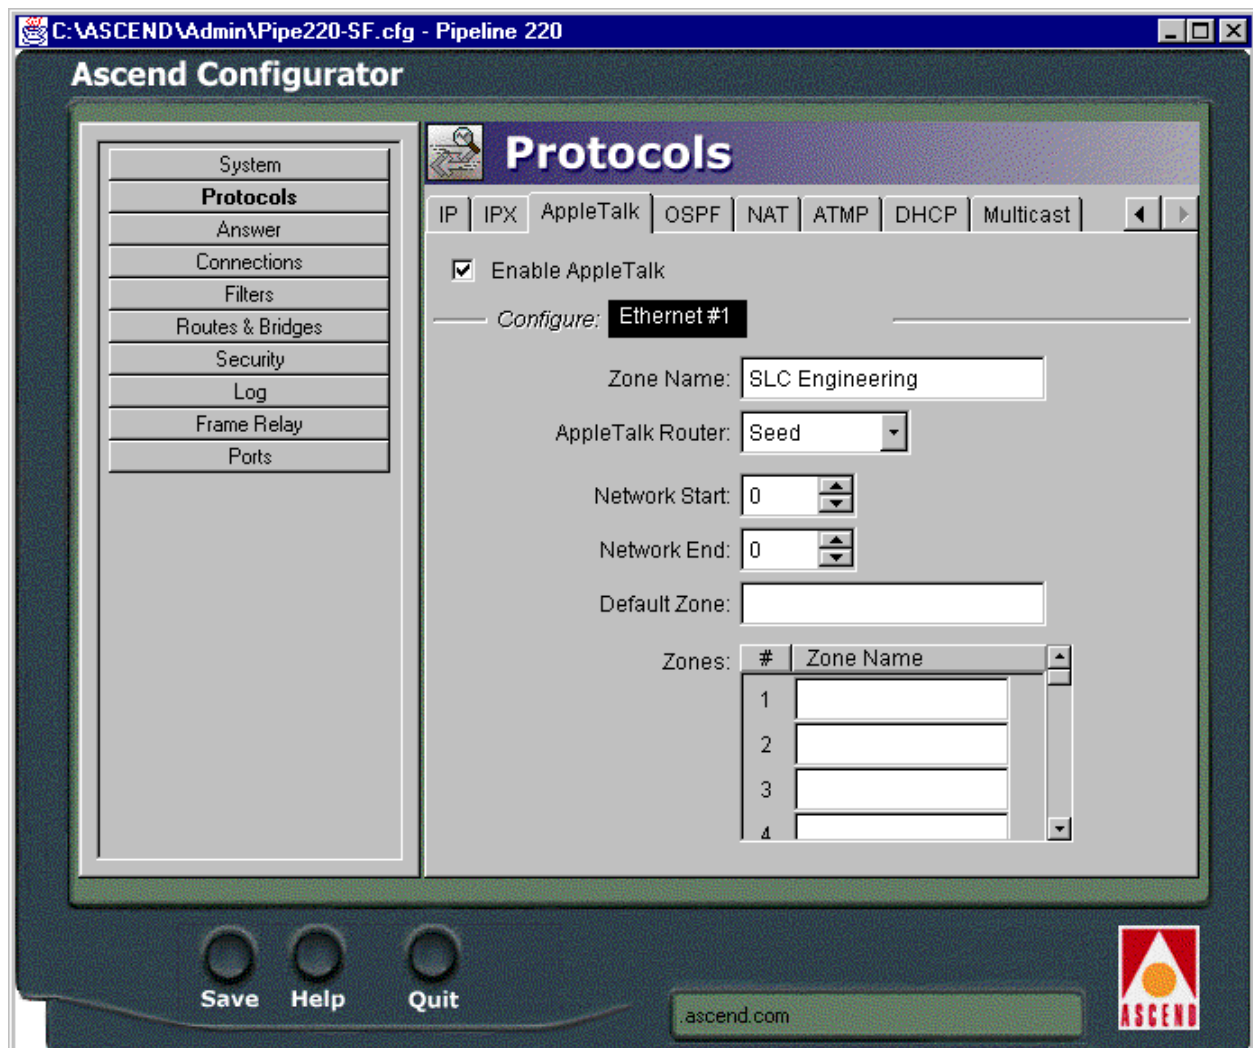
To configure AppleTalk routing, you must complete the steps outlined in “System-level AppleTalk routing parameters” and “Per-connection AppleTalk routing parameters” (if required).

System-level AppleTalk routing parameters

To set the system-level AppleTalk routing parameters:

- 1 On the left side of the Configurator, click the Protocols button.
- 2 On the right side of the Configurator, click the AppleTalk tab.

System-wide AppleTalk parameters appear on the right side of the Configurator:




- 3 In the upper-right section of the Configurator, select the Ethernet interface you want to configure, either Ethernet #1 or Ethernet #2.
Each Ethernet interface has its own unique set of system-wide parameters.
- 4 Set Zone Name to the zone in which the Pipeline 220 is located.

- 5 Set the AppleTalk Router parameter to specify whether the router is a seed or non-seed router.
A seed router has a manually defined network configuration. When a non-seed router boots, it has no local network configuration. It examines local network traffic and learns its local network configuration.
Note: You should configuring the Pipeline 220 as a non-seed router provided there is *at least one* seed router on the local network. Having only one seed router on a local network simplifies potential network configuration changes. Should you need to change the network numbering, only the seed router needs to be reconfigured. The remaining non-seed routers simply need to be rebooted to learn the changes.
- 6 If the Ascend unit is to be a seed router, specify the Network Start and Network End
If there are other seed routers sharing the Pipeline 220's network segment, this information must be identical on *all* routers that *share the network segment*. If there are no other seed routers, every network number from Network Start to Network End must be unique for the entire internet. Valid network numbers are of from 1 to 65,534.
- 7 Set Default Zone and any other zones assigned to the local AppleTalk network segment.
The Default Zone is assigned to any AppleTalk device that is connected to the Pipeline 220's local Ethernet segment and has not explicitly been assigned to another zone.
The Default Zone and additional zone list need to be identical for any AppleTalk router sharing the local network segment.
Note: Zones can be shared across network segments.
- 8 Enable the new settings, as described in "Saving the settings" on page 11-6.

Saving the settings

The Pipeline 220 settings you configure with the Ascend Configurator do not take effect until you upload them to the Pipeline 220. You use the Save command to upload the changes. Alternatively, you can save the setting to a text file for subsequent use, or to another device. To save the settings:

- 1 In the lower-left corner of the Configurator, click Save
A dialog box appears, prompting you for a save method:



- 2 Select one of the following:
 - If you are uploading the configuration to the same device, select “Upload changes to *<Pipeline 220 name>*.”
 - If you are saving the configuration to a text file on your computer, select “Save a copy under a new filename” and enter the filename to which the configuration should be saved.
 - If you are uploading the configuration to a different device, select “Upload to the Ascend product at:” and enter the IP address of the device.

- 3 Click Save.

If you selected “Upload changes to *<Pipeline 220 name>*” in step 2, the Ascend Configurator uploads the configuration file to the Pipeline 220 via TFTP. If you selected one of the other options, the Configurator sends the file to the location you specified. If you specified a remote device, transmission is via TFTP.

Per-connection AppleTalk routing parameters

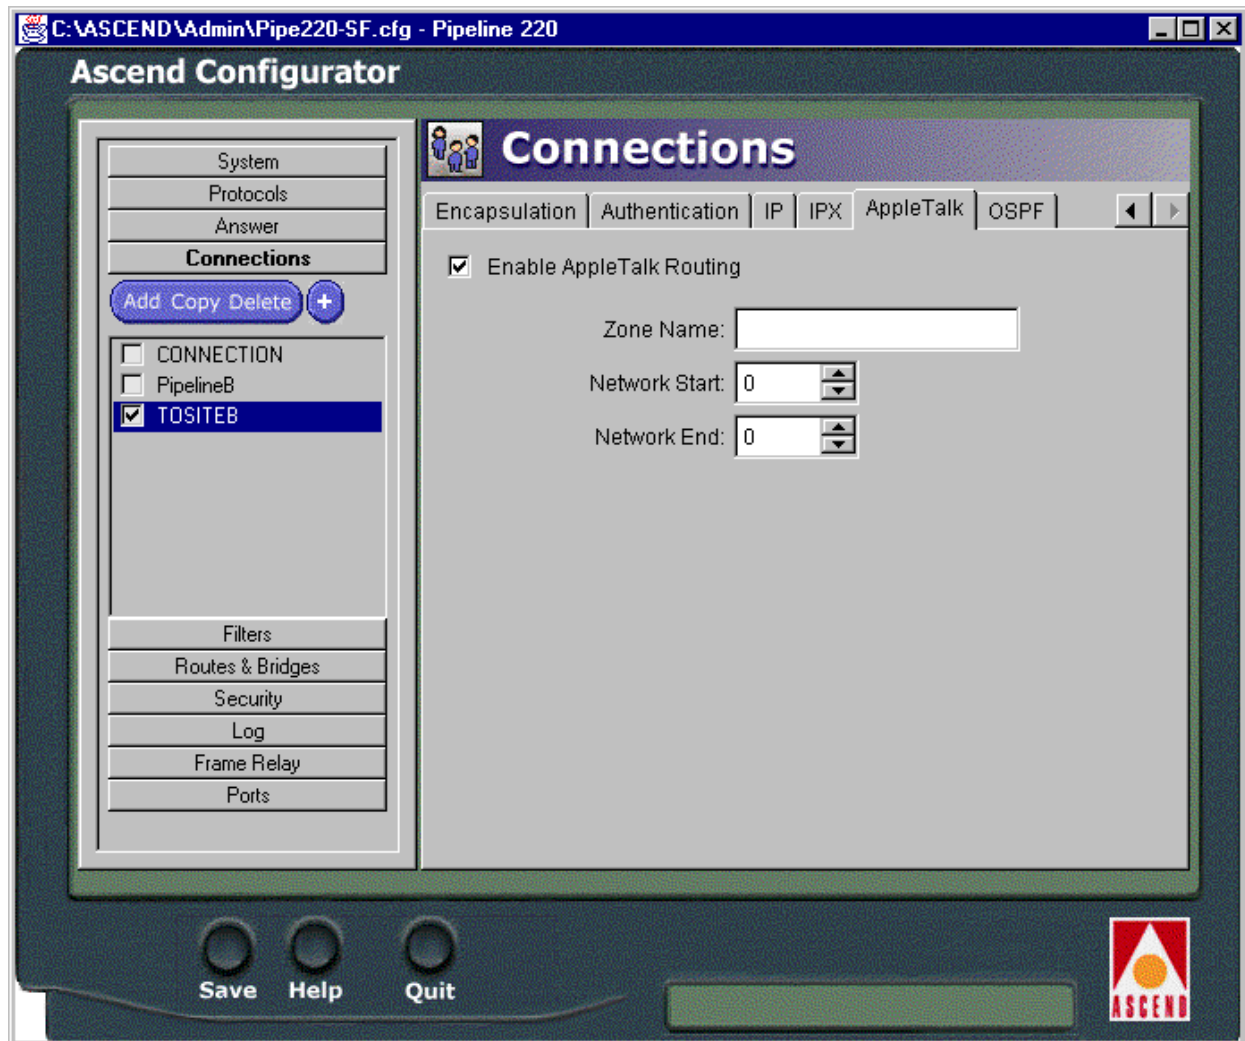
To set per-connection AppleTalk routing parameters:

- 1 On the left side of the Configurator, click the Connections button.
The Add/Copy/Delete window appears on the left side of the Configurator.

Configuring AppleTalk Routing

Configuring AppleTalk routing

- 2 In the Add/Copy/Delete window, highlight a connection.
The parameters for the connection you select appear on the right.



- 3 On the right side of the Configurator, click the Encapsulation tab.
- 4 Set Encapsulation to PPP.
The Pipeline 220 supports AppleTalk routing over PPP-encapsulated links.
- 5 Select the AppleTalk tab.
- 6 Select the Enable AppleTalk Routing check box.
- 7 Set Zone Name to the zone to which the remote AppleTalk router belongs.
- 8 Set Network Start and Network End to the network range to which the remote AppleTalk router belongs.
- 9 Enable the new settings, as described in “Saving the settings” on page 11-6.

Configuring Packet Bridging

This chapter covers the following topics:

Introduction to Ascend bridging	12-1
How the Pipeline establishes a bridging connection	12-3
Enabling bridging.	12-3
Managing the bridge table	12-7
Example of a bridged connection.	12-10

Introduction to Ascend bridging

This section provides an overview of packet bridging and explains how the Pipeline 220 brings up a bridging connection.

The Pipeline 220 is used as a bridge primarily to provide connectivity for protocols other than IP, IPX, and AppleTalk, although it can also be used for joining segments of an IP, IPX, or AppleTalk 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.

The most common uses of bridging in the Pipeline 220 are to:

- Provide any nonrouted protocol connectivity with another site
- Link any two sites so that their nodes appear to be on the same LAN
- Support protocols, such as BOOTP, that depend on broadcasts to function.

Disadvantages of bridging

Bridges examine *all* packets on the LAN (termed *promiscuous mode*), so they incur greater processor and memory overhead than routers. On heavily loaded networks, this increased overhead can result in slower performance.

Routers have other advantages over bridging. Because they examine packets at the network layer (instead of the link layer), you can filter on logical addresses, providing enhanced security and control. In addition, routers support multiple transmission paths to a given destination, enhancing the reliability and performance of packet delivery.

How a bridged WAN connection is initiated

When the Pipeline 220 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 local Ethernet segment (the segment to which the Pipeline 220 is connected).
- A broadcast address.

The important thing to remember about bridging connections is that they operate on physical and broadcast addresses, not on 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. For example:

0000D801CFF2

If the Pipeline 220 receives a packet whose destination MAC address is not on the local network, it first checks its internal bridge table (for a description of the table, see “Transparent bridging” on page 12-7). If it finds the packet's destination MAC address in its bridge table, the Pipeline 220 dials the connection and bridges the packet.

If the address is *not* specified in its bridge table, the Pipeline 220 checks for active sessions that have bridging enabled. If there are one or more active bridging links, the Pipeline 220 forwards the packet across *all* active sessions that have bridging enabled.

Broadcast addresses

A broadcast address is recognized by multiple nodes in 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. When configured as a router only, the Pipeline 220 discards broadcast packets. When configured as a bridge, it forwards packets with the broadcast destination address across all active sessions that have bridging enabled.

ARP broadcast packets that contain an IP address specified in the bridge table are a special case. For details, see “Static bridge table entries” on page 12-7.

How the Pipeline establishes a bridging connection

The Pipeline 220 uses station names and passwords to sync up a bridging connection, as shown in Figure 12-1.




Figure 12-1. Negotiating a bridge connection (PPP encapsulation)

The system name assigned to the Pipeline 220 in System button > Info tab > Name must *exactly* match the device name specified in the Connection profile on the remote bridge, including case changes. Similarly, the name assigned to the remote bridge must exactly match the name specified in the Connections button > General tab > Station Name parameter of that Connection profile, including case changes.

Note: The most common cause of trouble when initially setting up a PPP bridging connection is that the wrong name is specified for the Pipeline 220 or the remote device. Often case changes are not specified, or a dash, space, or underscore is not entered.

Enabling bridging

The Pipeline 220 has a system-wide bridging parameter that must be enabled for any bridging connection to work. The Bridging parameter directs the Pipeline 220 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. (Even if no packets are actually bridged, running in promiscuous mode incurs greater processor and memory overhead than the standard mode of operation for the Ethernet controller.)

Enabling bridging on the Ethernet interface

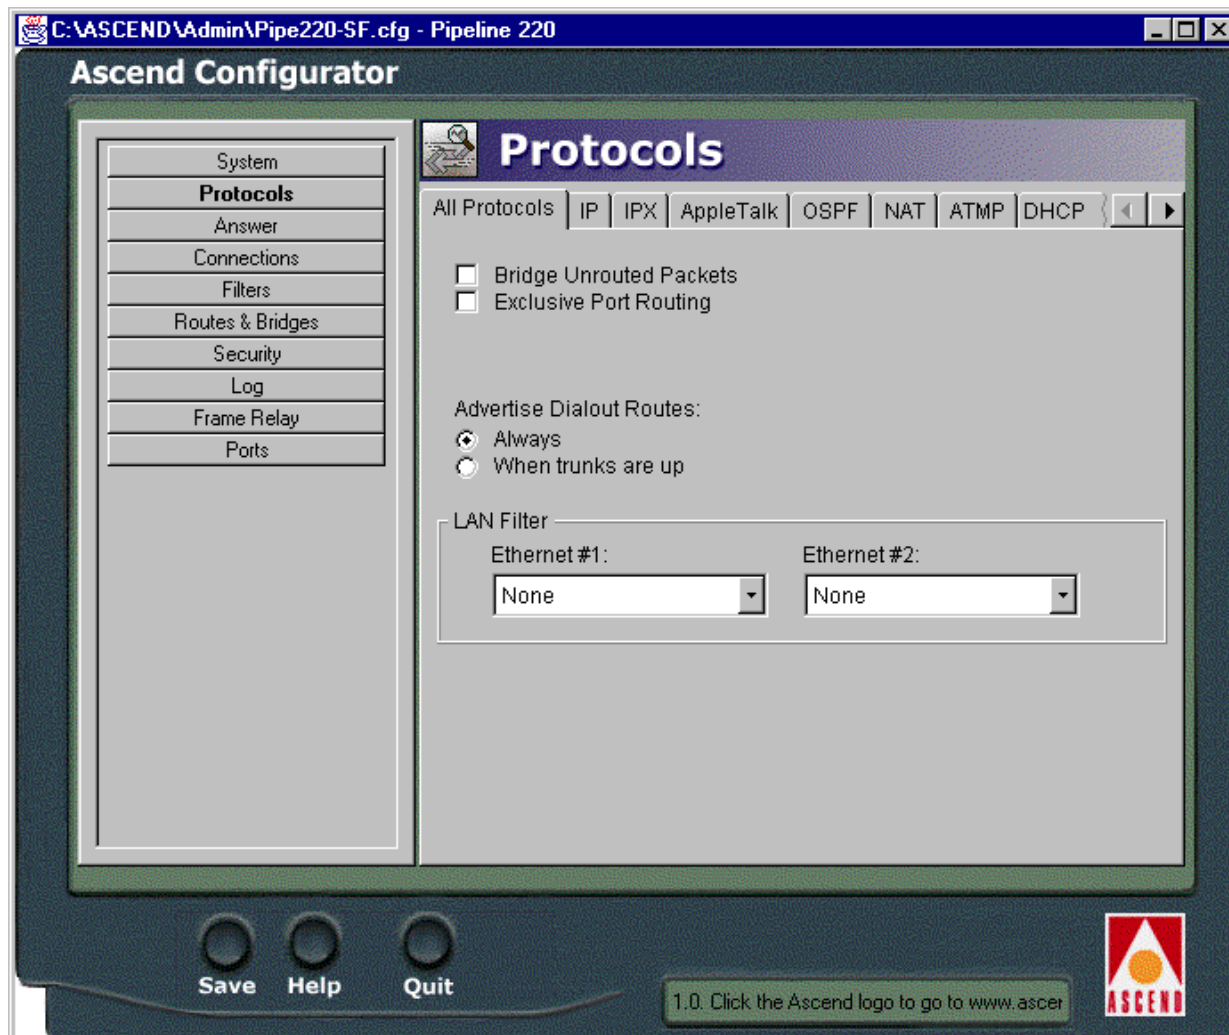
To bridge across a PPP-encapsulated link, you must enable bridging in the Answer profile.

To enable bridging on the Ethernet interface:

- 1 On the left side of the Configurator, click the Protocols button.

- 2 On the right side of the Configurator, click the All Protocols tab.

The following screen appears:



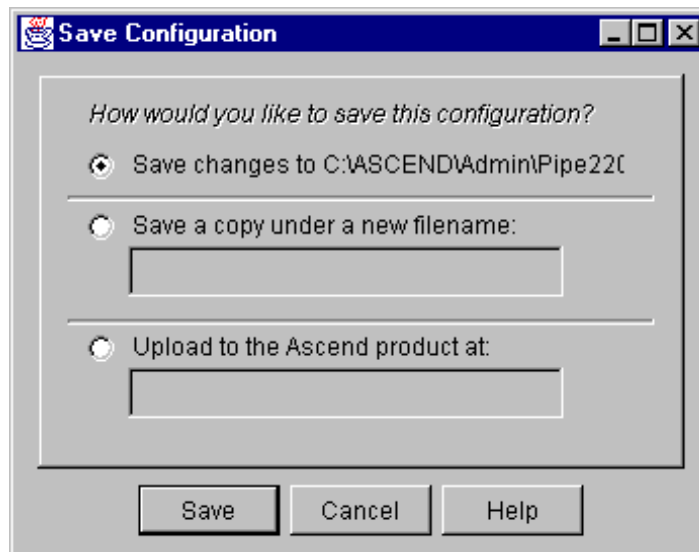
- 3 Select the Bridge Unrouted Packets check box.
- 4 Enable the new settings, as described in “Saving the settings” on page 12-4.

Saving the settings

The Pipeline 220 settings you configure with the Ascend Configurator do not take effect until you upload them to the Pipeline 220. You use the Save command to upload the changes. Alternatively, you can save the settings to a text file for subsequent use, or to another device. To save the settings:

- 1 In the lower-left corner of the Configurator, click Save.

A dialog box appears, prompting you for a save method:



- 2 Select one of the following:
 - If you are uploading the configuration to the same device, select “Upload changes to <Pipeline 220 name>.”
 - If you are saving the configuration to a text file on your computer, select “Save a copy under a new filename” and enter the filename to which the configuration should be saved.
 - If you are uploading the configuration to a different device, select “Upload to the Ascend product at:” and enter the IP address of the device.
- 3 Click Save.

If you selected “Upload changes to <Pipeline 220 name>” in step 2, the Ascend Configurator uploads the configuration file to the Pipeline 220 via TFTP. If you selected one of the other options, the Configurator sends the file to the location you specified. If you specified a remote device, transmission is via TFTP.

Bridging in the Answer Profile

Bridging must be enabled on both the local and remote side of a PPP-encapsulated link. Otherwise the link cannot bridge packets. In addition, PAP or CHAP authentication is required for unique identification of devices. (Specify PAP or CHAP in the Connection profiles for the links.)

Unlike an IP routing configuration, where the Pipeline 220 uniquely identifies the remote device by its IP address, a bridging configuration does not include a built-in way to identify incoming callers. For that reason, password authentication using PAP or CHAP is required unless IP routing is configured in each caller’s Connection profile.

To set Answer profile parameters for a bridging connection:

- 1 On the left side of the Configurator, click the Answer button.

Configuring Packet Bridging

Enabling bridging

- 2 On the right side of the Configurator, click the Routing tab.
Bridging and routing options appear on the right side of the Configurator:



- 3 Select the Enable Bridging check box. (If you cannot do so, you must first select the Protocols button > All Protocols tab > Bridge Unrouted Packets check box.)
- 4 Enable the new settings, as described in “Saving the settings” on page 12-4.

Managing the bridge table

To forward bridged packets to the correct destination network, the Pipeline 220 uses a bridge table that associates end nodes with particular connections. It builds this table dynamically (transparent bridging). 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 99 destination nodes and their connection information in Bridge profiles.

Transparent bridging

The Pipeline 220 is a transparent bridge (also termed a *learning bridge*). It keeps track of where a particular address is located, and of the Connection profile that specifies the interface to which the packet should be forwarded. As it forwards a packet, the Pipeline 220 logs the packet's source address and creates a bridge table that associates node addresses with a particular interface.

For example, Figure 12-2 shows the physical addresses of some nodes on the local Ethernet and at a remote site. The Pipeline 220 at site A is configured as a bridge.




Figure 12-2. How the Pipeline 220 creates a bridging table

The Pipeline 220 at site A gradually learns addresses on both networks by looking at each packet's source address, and it develops a bridge table like this:

0000D801CFF2	SITEA
080045CFA123	SITEA
08002B25CC11	SITEA
08009FA2A3CA	SITEB

Entries in the Pipeline 220 unit's bridge table must be relearned within a fixed aging limit, or they are removed from the table.

Static bridge table entries

You can specify up to 99 static bridge table entries for the Pipeline 220.

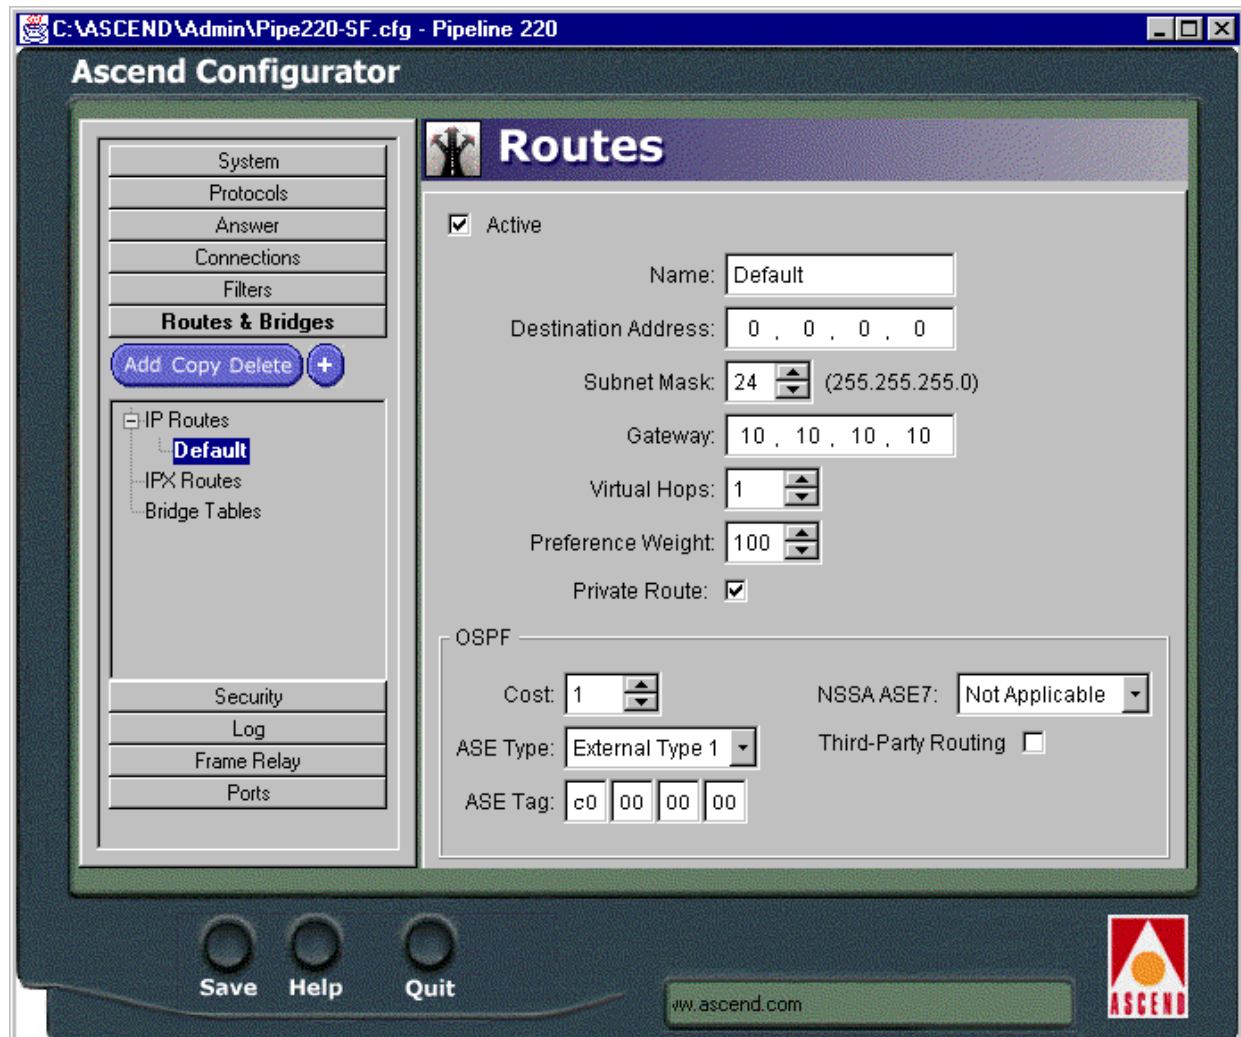
To define a static bridge table entry:

- 1 On the left side of the Configurator, click the Routes & Bridges button.

Configuring Packet Bridging


Managing the bridge table

Under the Routes & Bridges button, the Configurator displays the Add/Copy/Delete window:



- 2 In the Add/Copy/Delete window, click Bridge Tables.
- 3 Click Add.

The Configurator displays the New Bridge Table dialog box:



- 4 Enter the physical (MAC) address of the remote host. For example:

Ethernet Address=0080AD12CF9B

For a description of physical addresses, see “Physical addresses and the bridge table” on page 12-2. You must get this address from the administrator of the far-end device.

After you enter the address and click the OK button, the right side of the Configurator displays your new bridge table entry:



- 5 If the far-end is a segment of the local IP network, specify an address on that segment.
For example:
Network Address=10.1.2.133
- 6 Enable the new settings, as described in “Saving the settings” on page 12-4.

Example of a bridged connection

This section shows how to configure bridging for a Pipeline 220 connecting to a remote site. It describes a sample configuration that does not show the link-specific settings, or additional routing settings that might be appropriate at your site. It focuses only on bridging. (For details about each parameter, see the *Configurator Online Help*.)

In this example, two segments of an IP network are connected across the WAN, as shown in Figure 12-3. The two Pipeline 220 units are configured as bridges. A bridged connection at the link layer requires a bridge at both ends of the connection. (The most common cause of trouble when initially setting up a bridging connection is that the wrong name is specified for the Pipeline 220 or the remote device. Often case changes are not specified correctly, or a dash or underscore is entered incorrectly. Make sure you type the name exactly as it appears in the remote device.)

This example assumes that bridging has been enabled in both Pipelines by selecting the Protocols button > All Protocols tab > Bridge Unrouted Packets check box, as described in “Enabling bridging” on page 12-3. It also assumes you have enabled bridging by selecting the Answer button > Routing tab > Enable Bridging check box, as described in “Bridging in the Answer Profile” on page 12-5

In this example, two segments of an IP network are connected across the WAN, as shown in Figure 12-3.




Figure 12-3. Example of a bridged connection

Configuring the Site A Pipeline 220 for a bridged connection consists of assigning a system name to the Pipeline 220 unit, configuring a bridging connection, and defining a static bridge-table entry. The settings have to be saved after configuring the connection, and again after defining the table entry, because they do not take effect until the Pipeline 220 uploads them.

To avoid unnecessary traffic across the WAN, you can configure the Pipeline 220 to reply to ARP requests for any remote device for which it has a bridge-table entry.

To configure the site A Pipeline 220 for a bridged connection, you assign a name to the Pipeline 220, configure and save the bridging Connection profile, and configure and save a static bridge-table entry.

Assigning a name

Assign a name to the site A Pipeline 220:

- 1 On the left side of the Configurator, click the System button.
- 2 On the right side of the Configurator, click the Info tab.

- 3 Assign the Pipeline 220 a system name with the Name parameter.
Bridged connections use Name for authentication.

Configuring a bridging connection

To configure a bridging connection:

- 1 On the left side of the Configurator, click the Connections button.
- 2 On the right side of the Configurator, click the General tab.
- 3 Set the Station Name parameter. For example:
`Station=SITEBGW`
- 4 Click the Authentication tab.
- 5 Set Authentication to the type of authentication supported by local and remote devices.
For example:
`Authentication=CHAP`
- 6 Set Dial-Out Password to the password the Pipeline 220 sends to the remote device to authenticate itself.
- 7 Set Dial-In Password to the password the Pipeline 220 expects from the remote device.
- 8 On the upper-right section of the Configurator, click the IP tab.
- 9 Select the Enable IP Routing check box.
- 10 Enable the new settings, as described in “Saving the settings” on page 12-4.

Defining a static bridge-table entry

To define a static bridge table entry:

- 1 On the left side of the Configurator, click the Routes & Bridges button.
Under the Routes & Bridges button, the Configurator displays the Add/Copy/Delete window.
- 2 In the Add/Copy/Delete window, click Bridge Tables.
- 3 Click Add.
The Configurator displays the New Bridge Table dialog box.
- 4 Enter the physical (MAC) address of the remote host. For example:
`Ethernet Address=0080AD12CF9B`
For a description of physical addresses, see “Physical addresses and the bridge table” on page 12-2. You must get this address from the administrator of the far-end device.
After you enter the address, the right side of the Configurator displays your new Bridge Table entry.
- 5 If the far-end is a segment of the local IP network, specify an address on that segment.
For example:
`Net Address=10.2.3.100`
- 6 Enable the new settings, as described in step on page 1-4.

Configuring proxy mode on the Pipeline 220

If you are bridging between two segments of the same IP network, you can use the Net Address parameter in a Bridge profile to enable the Pipeline 220 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 220 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 220 acts as a proxy for the node that actually has that address.

Defining Static Filters

This chapter covers the following sections:

Introduction to Ascend filters	13-1
Overview of Filter profiles	13-6
Examples of filters	13-9

Introduction to Ascend filters

Ascend filters define packet conditions. 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.

How filters work

Without filtering, the Pipeline 220 forwards all packets. The conditions specified within a filter can specify not to forward certain packets, or not to forward *any* packets *except* those defined in the filter. The conditions also specify whether the Pipeline will examine inbound packets, outbound packets, or both. (For a more detailed discussion of specifying conditions, see “Overview of Filter profiles” on page 13-6.)

A filter’s forwarding action affects the actual data stream. Certain packets are dropped or forwarded, as specified in the filter conditions. Filters are often used for network security purposes, but they can be used for any purpose that requires the Pipeline to drop or forward only specific packets. For example, you can use filters to drop packets addressed to particular hosts or to prevent broadcasts from going across the WAN. On the other hand, you can use filters to allow only specific devices to be accessed by users across the WAN.

How to apply filters

You can apply a filter to the Answer profile, in which case it affects all WAN connections that do not have Connection profiles. For the connections that do, you can apply filters to their Connection profiles. You can also apply filters to the Ethernet interface. (Ascend units that accept switched calls can use *call* filters to prevent unnecessary connections. But they are not applicable to the Pipeline 220, which supports nailed connections only. Therefore, this chapter discusses only *data* filters.)

Applying a filter to the Answer profile

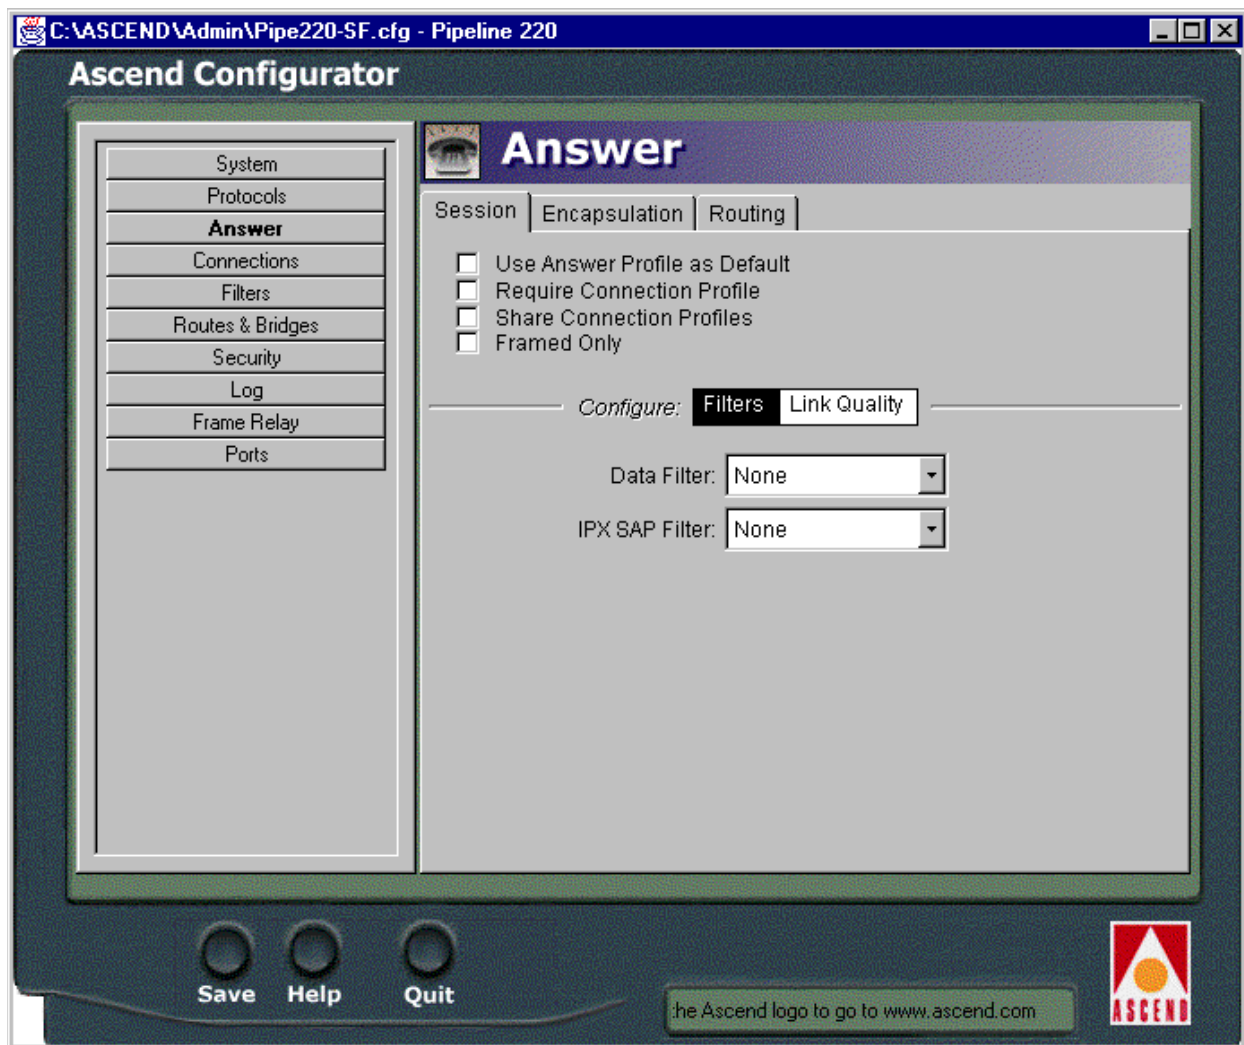
A filter applied to the Answer profile does not take effect if the WAN link uses a Connection profile. For links that do not use Connection profiles, the filter takes effect when the connection goes from an offline state to an active state.

Filters applied in the Answer profile are not used if the WAN link uses a Connection profile. When a link is brought up, any configured filter in the Answer button > Session tab > Filter button > Data Filter parameter is not used. If no filter is configured in the Connection button > General tab > Data Filter parameter, then no filter is applied.

After you have *defined* a filter, you can *apply* it to the Answer profile as follows:

- 1 On the left side of the Configurator, click the Answer button.
- 2 On the right side of the Configurator, click the Session tab.
- 3 On the right side of the Configurator, click the Filter button.

Filter parameters appear in the lower-right section of the Configurator:



- 4 Apply the filter. For example:

Data Filter=IP Spoof


No filter is applied if Data Filter is set to None (the default). To apply a filter, specify its Filter Name.

- 5 Save the settings, as described in “Saving the settings” on page 13-3.

Saving the settings

The Pipeline 220 settings you configure with the Ascend Configurator do not take effect until you upload them to the Pipeline 220. You use the Save command to upload the changes. Alternatively, you can save the setting to a text file for subsequent use, or to another device. To save the settings:

- 1 In the lower-left corner of the Configurator, click Save
A dialog box appears, prompting you for a save method:



- 2 Select one of the following:
 - If you are uploading the configuration to the same device, select “Upload changes to <Pipeline 220 name>.”
 - If you are saving the configuration to a text file on your computer, select “Save a copy under a new filename” and enter the filename to which the configuration should be saved.
 - If you are uploading the configuration to a different device, select “Upload to the Ascend product at:” and enter the IP address of the device.

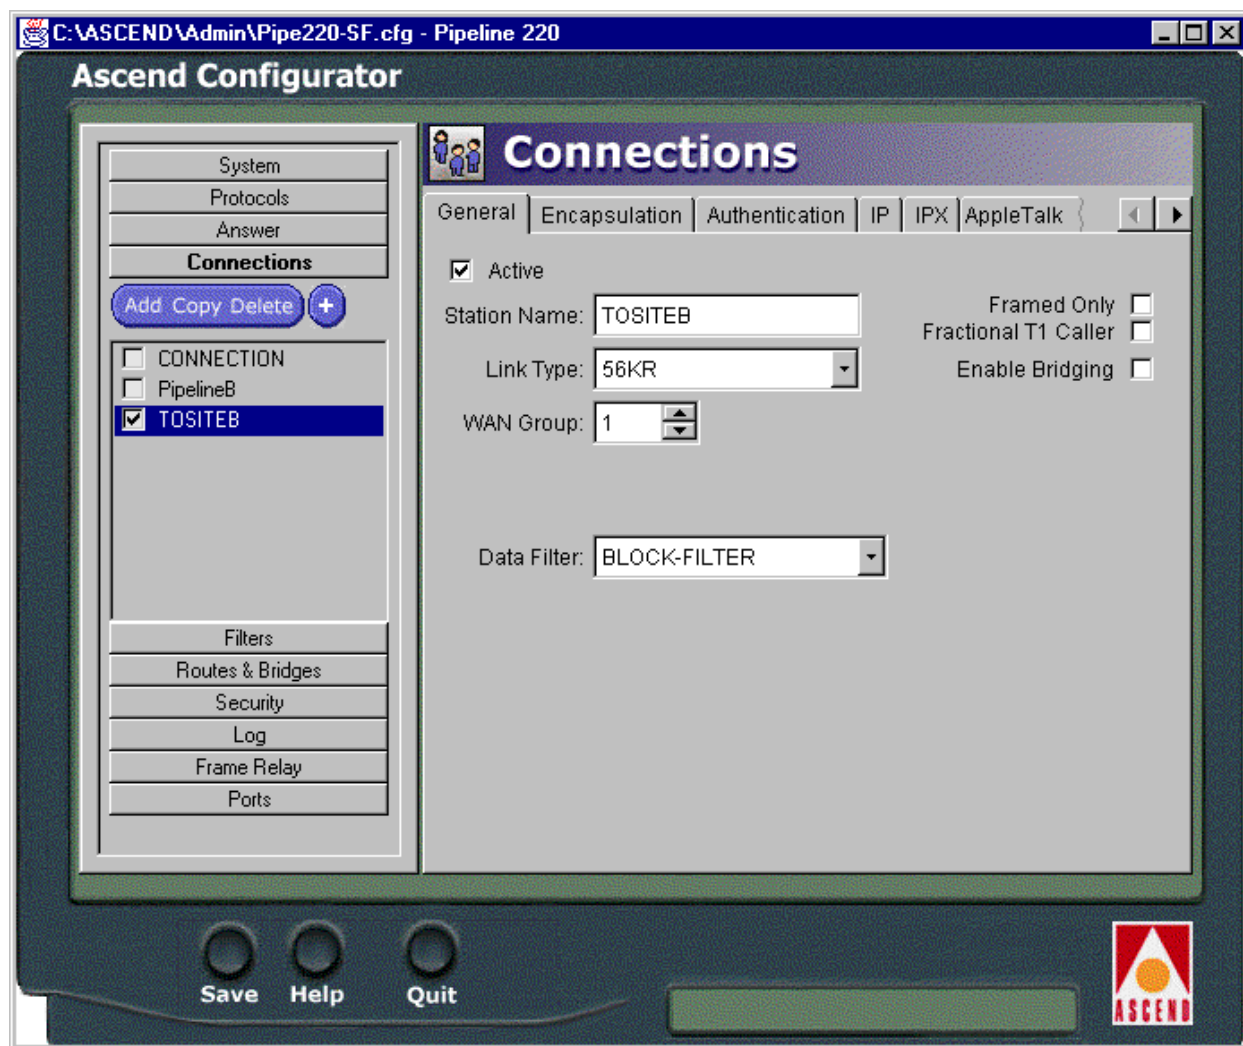
- 3 Click Save.

If you selected “Upload changes to <Pipeline 220 name>” in step 2, the Ascend Configurator uploads the configuration file to the Pipeline 220 via TFTP. If you selected one of the other options, the Configurator sends the file to the location you specified. If you specified a remote device, transmission is via TFTP.

Applying a filter to a Connection profile

After you have defined a filter, you can apply it to a specific Connection profile to specify which packets will be allowed to cross a WAN interface. Proceed as follows:

- 1 On the left side of the Configurator, click the Connections button.
If you set a filter in a Connection profile, it applies to that specific connection. (Unlike a filter in the Answer profile, which is applied to all unauthenticated connections.)
- 2 Access the Data Filter parameter.
For a Connection profile, the parameter is located in the Connections button > General tab:



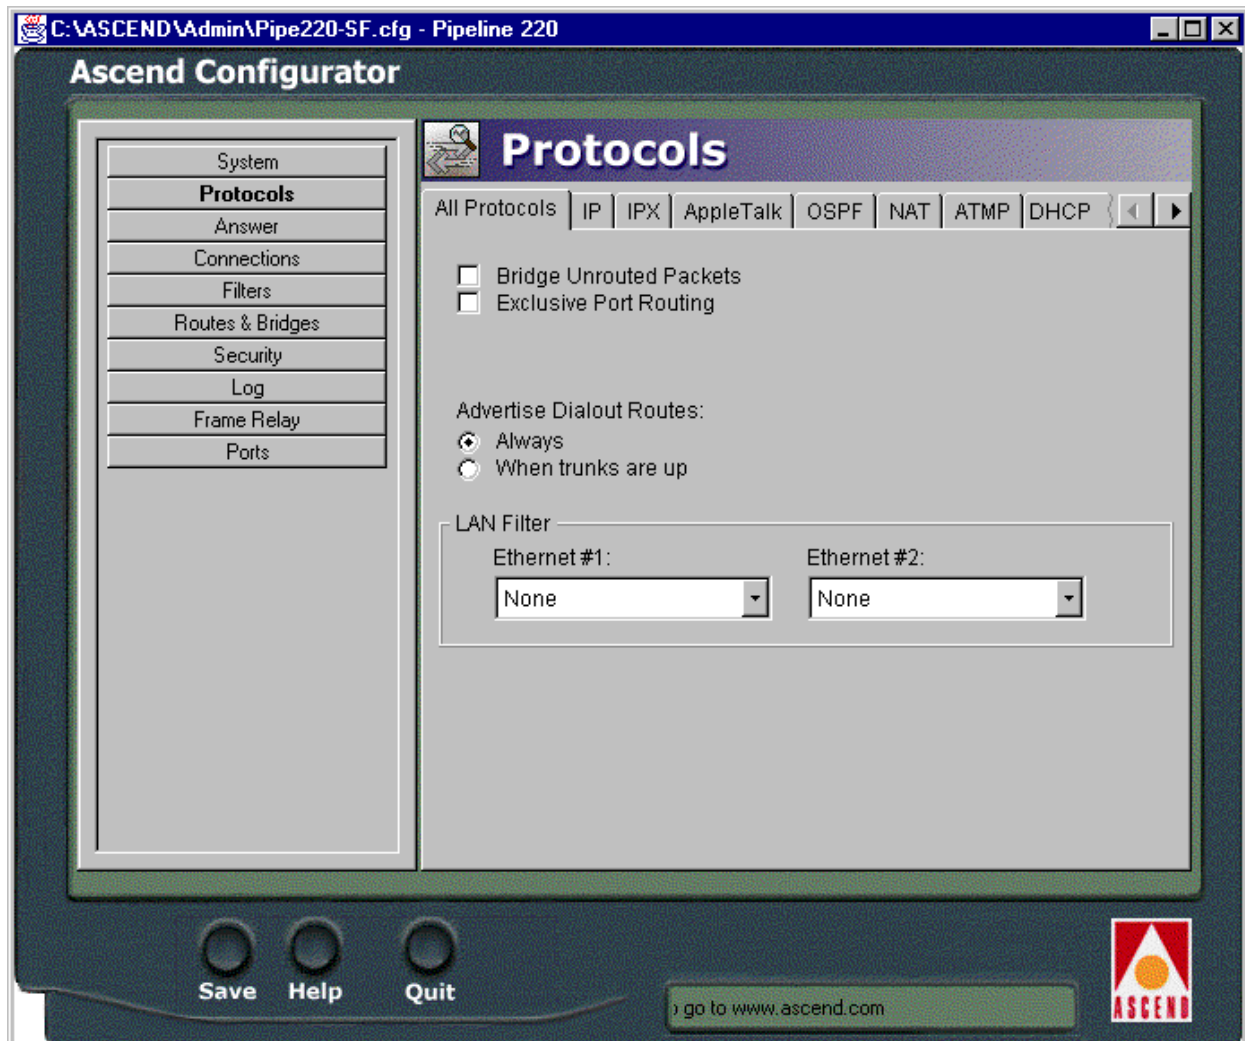
- 3 Save the settings, as described in “Saving the settings” on page 13-3.

Applying a filter to the Ethernet interface

After you have defined a filter, you can apply it as follows to specify which packets will be allowed to cross Ethernet interface:

- 1 On the left side of the Configurator, click the Protocols button.

- 2 On the right side of the Configurator, click the All Protocols tab.
LAN Filter parameters, one for each Ethernet interface, appear in the lower-right section of the Configurator:



- 3 Apply the filter.
For example:
`Data Filter=IP Spoof`
If this parameter is set to None (the default), no filter is applied. To apply a filter, specify its Filter Name.
- 4 Save the settings, as described in “Saving the settings” on page 13-3.

A filter applied to the Ethernet interface takes effect immediately. If you change the Filter Profile definition, the new filters are applied as soon as you save the Filter Profile, if you upload the changes to the Pipeline 220.

Overview of Filter profiles

The three basic types of filters are Generic filters, IP filters, and IPX filters. Generic filters examine the byte- or bit-level contents of packets. They focus on bytes or bits at particular locations, and compare the contents of a location with a value defined in the filters. Protocol specifications are usually the best source of the information you need for effective use of Generic filters.

IP filters examine higher-level fields specific to IP, TCP, and UDP packets. IP filters focus on known fields in IP packets, such as source or destination address, protocol number, and so forth. They operate on logical information, which is relatively easy to obtain.

Figure 13-1 shows how filters are organized and the terminology used to describe each part of a filter.




Figure 13-1. Filter terminology

The menus shown in Figure 13-1 are nested submenus below the Filters menu. The following table describes the structure of Ascend filters:

Menu	Description
Filters menu	The Configurator shows the Filters menu, a list of currently defined filters, in the Add/Copy/Delete window when you click the Filters button. When applying a filter, you identify it by its Filter Name.
Filter profile	Each Filter profile contains a set of filter conditions for a particular Filter Name. When you highlight a Filter Name in the Add/Copy/Delete window, its Filter profile is available on the right side of the Configurator.
Input or Output filters	At the top level of a Filter profile are submenus labeled Input Filters and Output Filters. Each submenu contains a list of 12 filters. The conditions you define within those <i>Input Filters</i> or <i>Output Filters</i> (or both) are applied to the inbound or outbound packet stream, respectively, in the order in which they appear (1–12). For details, see “Filtering inbound and outbound packets” on page 13-7.
Generic, IP, or IPX filters	Each <i>Input Filter</i> and <i>Output Filter</i> can be of type GENERIC, IP, or IPX. Once you assign a type, you can open the corresponding submenu to define the packet-level filter conditions. For details, see “Selecting filter type, activating the filter, and defining the conditions” on page 13-7.

Menu	Description
Filter conditions	Filter conditions specify the actual packet characteristics that will be examined in the data stream. Generic filters conditions specify locations and values that can be found within any packet. IP filter conditions specify IP-specific packet characteristics, such as address, mask, and port. IPX filter conditions specify IPX-specific packet characteristics, such as socket, network number, and network node.

Filtering inbound and outbound packets

At the top level of a Filter profile, you assign a name and click either the Input Filters or Output Filters tab.

The Pipeline 220 compares Input filters to received packets. If the filter is applied to the Ethernet interface, any Input filters are compared to packets from the Ethernet *into* the Pipeline. If applied as a filter on a WAN interface (applied to a Connection Profile), it affects packets from that WAN interface *into* the Pipeline.

The Pipeline 220 compares Output filters to packets to be sent. If the filter is applied to the Ethernet interface, any Output filters are compared to packets in the Pipeline 220 destined for the Ethernet interface. If applied as a filter on a WAN interface (applied to a Connection Profile), it affects packets in the Pipeline 220 destined for the WAN interface.

You can specify up to 12 Input Filters and 12 Output Filters in a Filter profile. They are applied in order from 1 to 12.

For security reasons, if a packet does not match *any* of the defined conditions, it is discarded.

However, if *only* Input filters are defined, the default action for *Output* filters is to forward all packets. The same is true in the other direction. If you define only Output filters, the default action for inbound packets is to forward them.

Selecting filter type, activating the filter, and defining the conditions

The Input Filters and Output Filters you define are applied to a packet in numerical order from one to twelve, provided that each filter has the Enabled check box selected. Clearing the Enabled check box for a filter prevents it from being applied.

After you open an Input Filter or an Output Filter, check the Enabled check box. Then, select the button corresponding to the type of filter conditions to be defined (Generic, IP, or IPX).

Generic filter conditions define bits and bytes within a packet. They are applied to all packet types. IP filter conditions apply only to TCP, IP, and UDP packets. IPX filter conditions apply only to IPX packets.

Defining generic filter conditions

If you select the Generic Packets button, you define conditions for a Generic filter. To define a Generic filter, you set the following parameters (for complete information about each parameter, see the *Configurator Online Help*):

Parameter	Description
Filter Action	Determines whether the Pipeline will forward a packet if it matches the definition or discard the packet if it matches.
Offset, Length, Mask, and Value	You use the Offset, Length, Mask, and Value parameters to define the exact location of certain bytes within a packet and the values of those bytes.
Comparison	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 Comparison is set to Comparison matches (the default), the filter is applied if the packet data are identical to the specified value. If Comparison is set to Does not match, the filter is applied if the packet data are not identical.
Link this condition to the next...	Specifies whether the current filter is linked to the one immediately following it. If you select this check box, the filter can examine multiple non-contiguous bytes within a packet by <i>linking</i> the current filter to the next one, so that the next filter is applied before the Filter Action decision is made. The match occurs only if <i>both</i> non-contiguous bytes contain the specified values. If the check box is cleared, the Filter Action decision is based on whether the packet matches the definition in this one filter.

Defining IP filter conditions

If you select the IP button, you define conditions related only to TCP, IP, and UDP packet filtering. These packets are compared regardless of whether the Pipeline 220 routes or bridges IP. An IP filter examines source addresses, destination addresses, IP protocol type and port, or a combination of these.

To define an IP filter, you set the following parameters (for complete information about each parameter, see the *Configurator Online Help*):

Parameter	Description
Filter Action	Determines whether the Pipeline will forward a packet if it matches the definition or discard the packet if it matches.
Source and destination address and mask	Specify the contents of the source or destination fields in a packet. Use the address parameter to specify the source or destination address, and the Mask parameter to mask out portions of the address (for example, to mask out the host number).

Parameter	Description
Protocol	Identifies 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, <i>Assigned Numbers</i> , by Reynolds, J. and Postel, J., October 1994. 1 — ICMP 5 — STREAM 8 — EGP 6 — TCP 9 — Any private interior gateway protocol (such as 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
Source and destination Port Compare and Port #	Specify whether to compare the protocol ports, which identify the application running over TCP/IP. The comparison can match a protocol port number that is less-than, greater-than, equal, or not-equal.
Match only established TCP connections	Set this parameter to compare a packet only if it is part of a TCP session that is already established.

Examples of filters

This section provides step-by-step examples that show how to specify Generic and IP filter conditions.

Example of a generic filter to handle AppleTalk broadcasts

This section shows how to define a Generic filter whose purpose is to prevent local AppleTalk Echo Protocol (AEP) and Name Binding Protocol (NBP) traffic from going across the WAN. The 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)

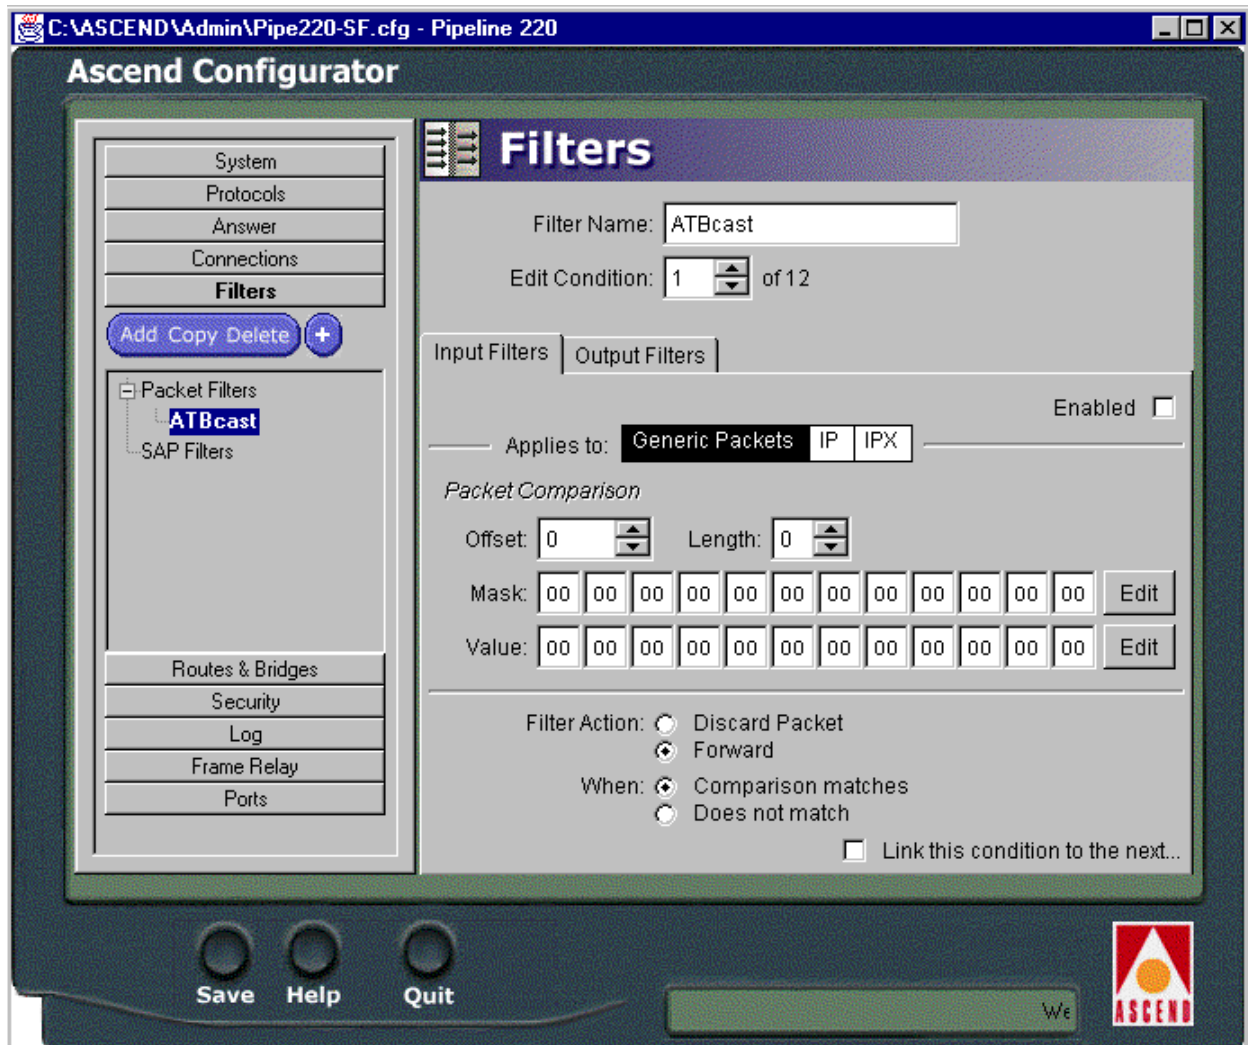
Setting up the filter

To set up the filter:

- 1 On the left side of the Configurator, click the Filters button.
The Add/Copy/Delete window appears in the lower-left section of the Configurator.
- 2 On the Add/Copy/Delete menu, click Add.
A Dialog box appears and prompts you for a value to apply as Filter Name:



- 3 Enter a name for the filter and select OK.
The right side of the Configurator displays the new filter. It has the name you entered, with default values for all other parameters:



- 4 Set Edit Condition to 1.
- 5 On the right side of the Configurator, click the Output Filters tab.
- 6 Select the Enabled check box on the right side of the Configurator.
- 7 Click the Generic Packets button.

Blocking ARP packets

Next, configure a filter condition that defines a location, within a packet, that indicates whether it is an ARP packet:

- 1 Set Offset=14
This filter condition is applied to every outgoing Ethernet frame. It causes the comparison to start 14 bytes into the Ethernet frame instead of at the beginning.
- 2 Set Length=8
The filter condition compares an eight-byte section of every Ethernet frame.
- 3 Set Mask=ffffffffffffffff
This Mask specifies that every bit of the 8 bytes will be compared.

Defining Static Filters

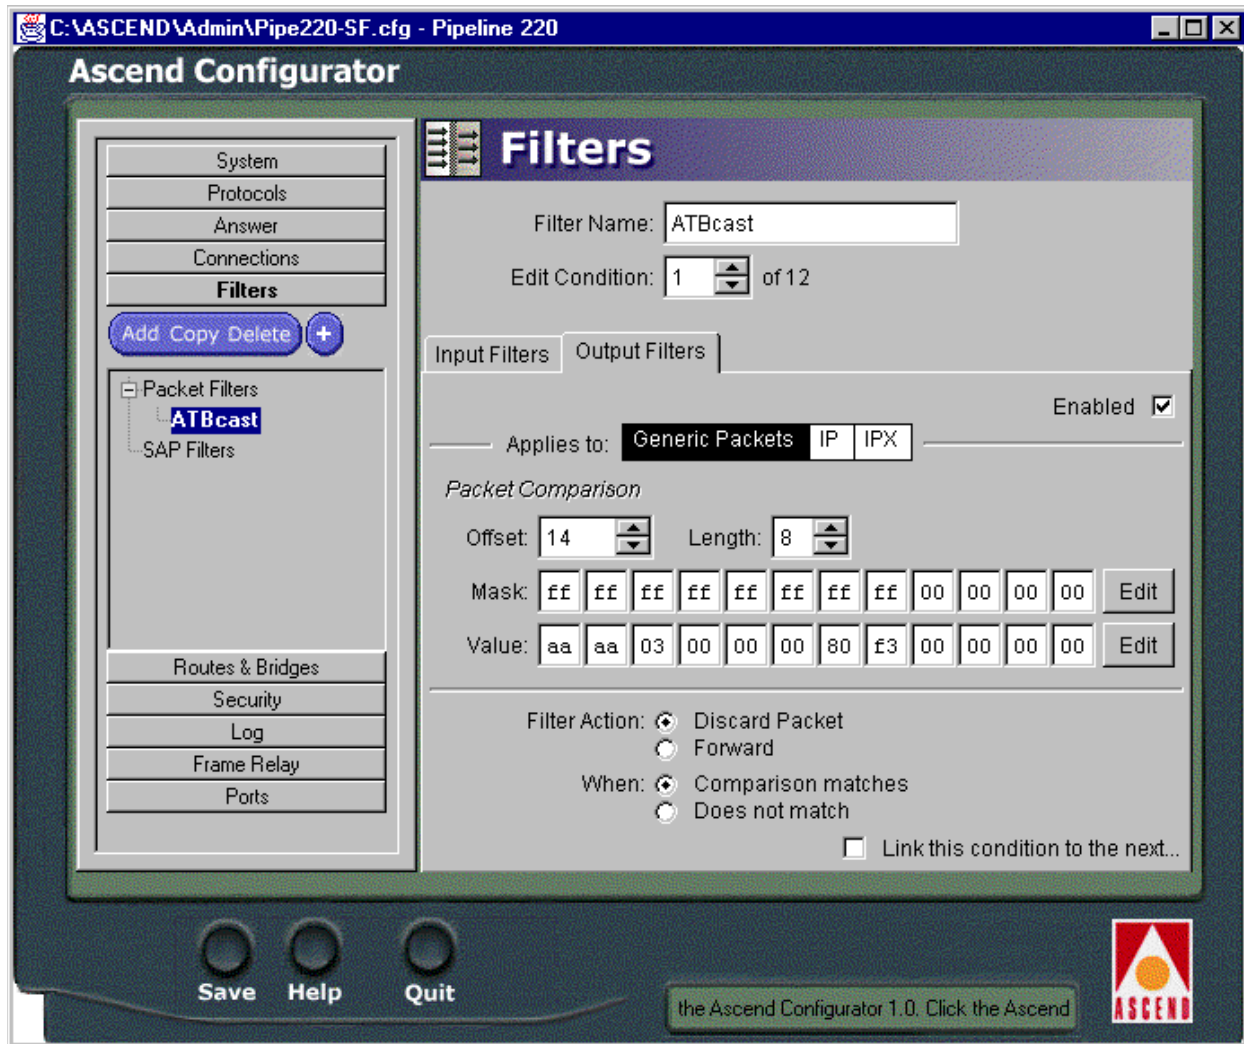
Examples of filters

In some cases, you might want to consider specific bits. Set the Mask parameter with a hexadecimal number that specifies zeroes for the bits that should not be compared. In this case, the last four bytes of the data being compared identify the Ethernet protocol type, and ARP is defined as Ethernet Protocol Type 0x80f3. Proceed as follows:

- 4 Set Value=aaaa0300000080f3
- 5 Set Filter Action to Discard Packet .
- 6 Set When to Comparison matches.

Steps 5 and 6 ensure that the Pipeline 220 will discard any ARP packet.

- 7 Clear the “Link this condition to the next...” check box.
Before continuing, verify the settings of the first sample filter condition:



Allowing non-AppleTalk traffic to pass

AppleTalk is Ethernet Protocol Type 0x809b, so to allow all non-AppleTalk traffic to pass through the Pipeline 220, define a filter that forwards any packet that does not have Type 0x809b. Proceed as follows:

- 1 Set Edit Condition to 2.
- 2 Select the Enabled check box on the right side of the Configurator.
- 3 Set Offset=14
This filter condition is applied to every outgoing Ethernet frame. It causes the comparison to start 14 bytes into the Ethernet frame instead of at the beginning.
- 4 Set Length=8
The filter condition compares an eight-byte section of every Ethernet frame.
- 5 Set Mask=ffffffffffffffff

Defining Static Filters

Examples of filters

This Mask specifies that every bit of the 8 bytes will be compared.

- 6 Set Value=aaaa03080007809b

The value 0x809b indicates that the packet in the Ethernet frame is an AppleTalk packet.

- 7 Set Filter Action to Forward.

- 8 Set When to Does not match.

Steps 7 and 8 ensure that the Pipeline 220 will forward any non-AppleTalk packet.

- 9 Clear the “Link this condition to the next...” check box.

Allowing AppleTalk AEP packets to pass

To configure a filter condition that allows AppleTalk AEP packets to pass through the Pipeline 220:

- 1 Set Edit Condition to 3.

- 2 Select the Enabled check box on the right side of the Configurator.

- 3 Set Offset=32

This filter condition is applied to every outgoing Ethernet frame. It causes the comparison to start 32 bytes into the Ethernet frame instead of at the beginning.

- 4 Set Length=3

The filter condition compares a three-byte section of every Ethernet frame.

- 5 Set Mask=ffffff0000000000

This Mask specifies that every bit of the 3 bytes will be compared.

- 6 Set Value=0404040000000000.

- 7 Set Filter Action to Forward.

Steps 7 and 8 ensure that any AEP packet will be forwarded by the Pipeline 220.

- 8 Set When to Comparison matches.

- 9 Clear the “Link this condition to the next...” check box.

Blocking broadcast AppleTalk packets

To configure a filter condition that does not allow broadcast AppleTalk packets to pass through the Pipeline 220, take the inverse approach, allowing non-broadcast AppleTalk traffic to pass:

- 1 Set Edit Condition to 4.

- 2 Select the Enabled check box on the right side of the Configurator.

- 3 Set Offset=32

This filter condition is applied to every outgoing Ethernet frame. It causes the comparison to start 32 bytes into the Ethernet frame instead of at the beginning.

- 4 Set Length=6

The filter condition compares a six-byte section of every Ethernet frame.

- 5 Set Mask=ffffffffffff0000

This Mask indicates that every bit of the 6 bytes will be compared.

- 6 Set Value=090007ffffffff000

AppleTalk broadcast traffic is identified by its use of a multicast address. This value specifies a packet that uses a multicast address.

- 7 Set Filter Action to Forward.

Steps 7 and 8 ensure that the Pipeline 220 will discard any broadcast AppleTalk packet.

- 8 Set When to Does not match.
- 9 Clear the “Link this condition to the next...” check box.

Blocking NBP packets

To block Name Binding Protocol (NBP) broadcast packets, define two linked filter conditions. These linked output filter conditions specify NBP lookup packets that use a wildcard device name. AppleTalk devices use NBP lookups to search for specific devices. For example, when you open a Macintosh Chooser to connect to a printer, the Macintosh sends NBP lookup packets that are responded to by any available printers. To define the first condition and set up the link:

- 1 Set Edit Condition to 5.
- 2 Select the Enabled check box on the right side of the Configurator.
- 3 Set Offset=32
This filter condition is applied to every outgoing Ethernet frame. It causes the comparison to start 32 bytes into the Ethernet frame instead of at the beginning.
- 4 Set Length=4
The filter condition compares a four-byte section of every Ethernet frame.
- 5 Set Mask=ff00ff0000000000
This Mask indicates that not every bit of the four bytes will be compared.
- 6 Set Value=0200022000000000
- 7 Set Filter Action to Forward.
- 8 Set When to Comparison matches.
- 9 Select “Link this condition to the next...” check box.
Selecting “Link this condition to the next...” in a filter condition links it with the next condition you configure. A packet must match conditions from *both* filter conditions to be forwarded. If a packet satisfies only one condition, the packet is discarded.

To define the second condition:

- 1 Set Edit Condition to 6.
- 2 Select the Enabled check box on the right side of the Configurator.
- 3 Set Offset=42
This filter condition is compared to every outgoing Ethernet frame. It causes the comparison to start 42 bytes into the Ethernet frame instead of at the beginning.
- 4 Set Length=2
The filter condition compares a two-byte section of every Ethernet frame.
- 5 Set Mask=ffff000000000000
This Mask indicates that every bit of the two bytes will be compared.
- 6 Set Value=013d000000000000
- 7 Set Filter Action to Forward.
- 8 Set When to Comparison matches.
- 9 Clear the “Link this condition to the next...” check box.

Discarding unrecognized packets

To define the final condition, specify that if a packet has passed through the previous filter conditions and has not matched any of them, it will be discarded:

- 1 Set Edit Condition to 7.
- 2 Select the Enabled check box on the right side of the Configurator.
- 3 Set Offset=0
This filter condition is applied to every outgoing Ethernet frame. It causes the comparison to start at the beginning of the frame.
- 4 Set Length=0
The filter condition compares nothing, so every Ethernet frame matches.
- 5 Set Mask=0000000000000000
- 6 Set Value=0000000000000000
- 7 Set Filter Action to Discard Packet.
Steps 7 and 8 ensure that the Pipeline 220 discards any unknown traffic.
- 8 Set When to Comparison matches.
- 9 Clear the “Link this condition to the next...” check box.

Saving the filter

To actually create the filter you have defined, see “Saving the settings” on page 13-3.

Example of an IP filter to prevent address spoofing

This section shows how to define an IP filter to prevent *spoofing* of local IP addresses. Spoofing IP addresses (not to be confused with DHCP spoofing described in Chapter 7, “IP Address Management”) is a technique whereby outside users pretend to be from the local network in order to obtain unauthorized access to the local network.

The conditions define Input filters that drop packets whose source address either is on the local IP network or is the loopback address (127.0.0.0). All other incoming packets are forwarded.

The conditions also define an Output filter that forwards every outbound packet that has a source address on the local network. All outbound packets with a nonlocal source address are discarded.

Note: This example assumes a local IP network address of 192.100.50.128, with a subnet mask of 255.255.255.192. Use your own local IP address and mask when defining a filter to prevent address spoofing.

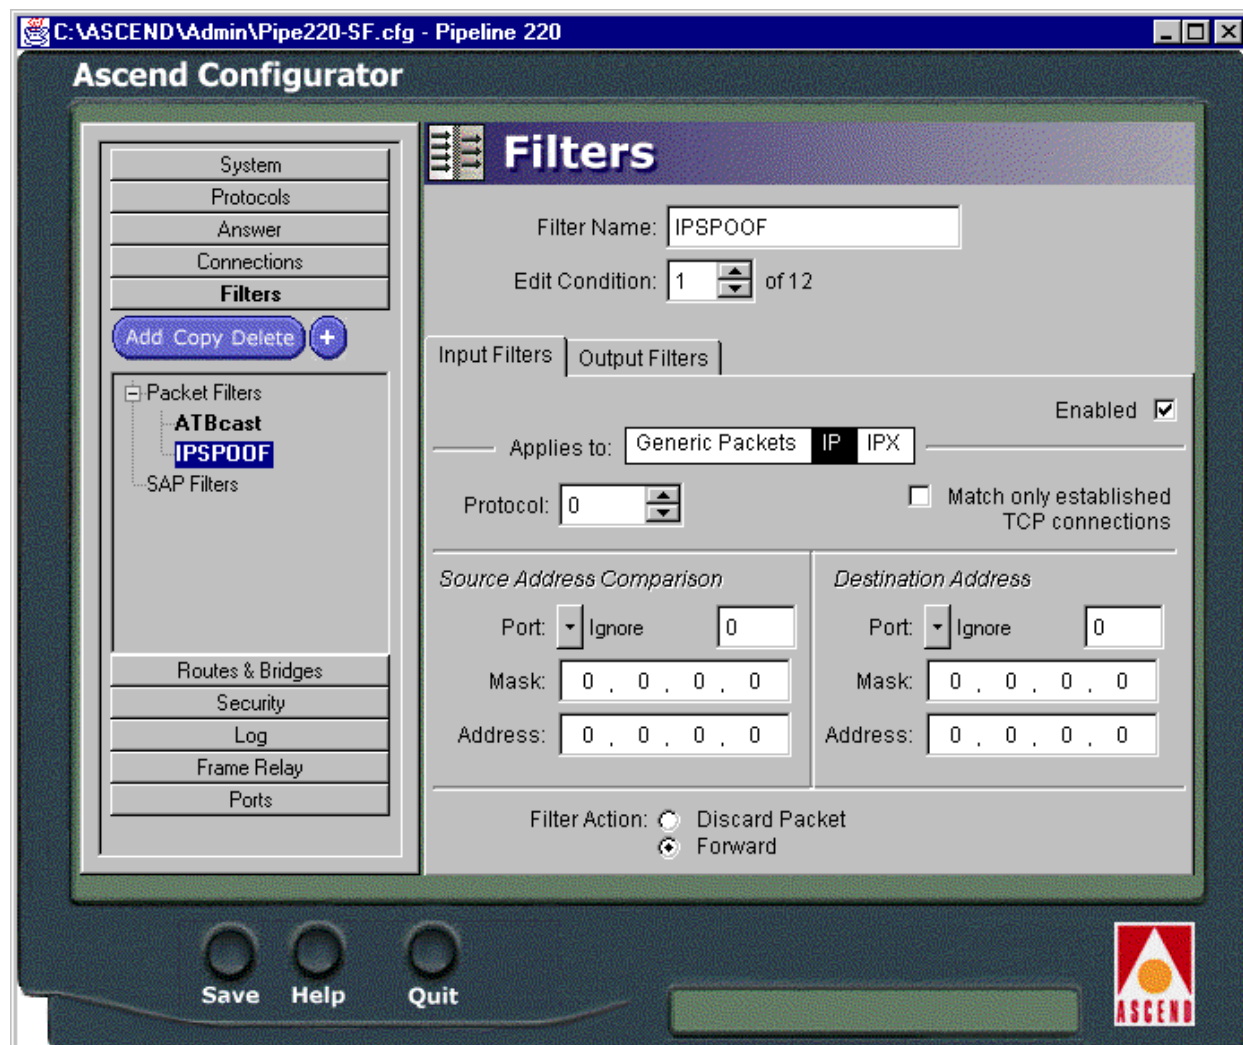
To define this IP filter, you first set it up, then define conditions to:

- Discard any incoming packet that has the local subnet as its source address.
- Discard any incoming packet that has the loopback address as its source address.
- Forward all incoming packets that have nonlocal source addresses.
- Forward all outbound packets that have local source addresses.

Setting up the filter

To begin defining this filter, proceed as follows:

- 1** On the left side of the Configurator, click the Filters button.
The Add/Copy/Delete window now appears in the lower-left section of the Configurator.
- 2** On the Add/Copy/Delete menu, click Add.
A dialog box appears and prompts you for a value to apply as Filter Name.
- 3** Enter a name for the filter and select OK. For example:
Filter Name=IPSpooF
- 4** Set Edit Condition to 1.
- 5** Click the Input Filters tab.
- 6** Select the Enabled check box on the right side of the Configurator.
- 7** Click the IP Packets button.
IP-filter parameters appear in the lower-right section of the Configurator:



Discarding incoming packets with the local subnet as the source address

Configure the first condition to specify the local subnet mask and IP address, and discard any incoming packet has the specified local address. Proceed as follows:

- 1 Set Protocol=0
- 2 Clear the “Match only established TCP connections” check box.
- 3 Set Port (Source)=Ignore
Now the Pipeline 220 will not compare the Port field in the source-address portion of the packet.
- 4 Set Mask (Source)=255.255.255.192
The first twenty six bits of the IP address (the first three octets plus two bits from the fourth) indicate the subnet. The remaining six bits indicate the host portion of the address. The Pipeline 220 will compare only the subnet portion of the source address.
- 5 Set Address (Source)=192.100.50.128
- 6 Set Port (Destination)=Ignore

- 7 Set Mask (Destination)=0.0.0.0
- 8 Set Address (Destination)=0.0.0.0
- 9 Set Filter Action to Discard Packet.

Discarding packets with a loopback source address

Configure Input condition two to discard any incoming packet specifying the loopback address as its source address:

- 1 Set Edit Condition to 2.
- 2 Select the Enabled check box on the right side of the Configurator.
- 3 Set Protocol=0
- 4 Clear the “Match only established TCP connections” check box.
- 5 Set Port (Source)=Ignore
Now the Pipeline 220 will not compare the Port field in the source-address portion of the packet.
- 6 Set Mask (Source)=255.0.0.0
The Pipeline 220 compares the first eight bits of the source address.
- 7 Set Address (Source)=127.0.0.0
- 8 Set Port (Destination)=Ignore
- 9 Set Mask (Destination)=0.0.0.0
- 10 Set Address (Destination)=0.0.0.0
- 11 Set Filter Action to Discard Packet.

Forwarding incoming packets that have nonlocal source addresses

Configure Input condition three to forward *any* incoming packet (designated as 0.0.0.0) specifying a non-local source address:

- 1 Set Edit Condition to 3.
- 2 Select the Enabled check box on the right side of the Configurator.
- 3 Set Protocol=0
- 4 Clear the “Match only established TCP connections” check box.
- 5 Set Port=Ignore
- 6 Set Mask (Source)=0.0.0.0
- 7 Set Address (Source)=0.0.0.0
- 8 Set Port (Destination)=Ignore
- 9 Set Mask (Destination)=0.0.0.0
- 10 Set Address (Destination)=0.0.0.0
- 11 Set Filter Action to Forward.

Forwarding outbound packets that have local source addresses

To complete the filter definition, configure an Output Filter to forward any outbound packet with a local source address:

Defining Static Filters

Examples of filters

- 1 Set Edit Condition to 1.
- 2 Click the IP button.
- 3 Select the Enabled check box on the right side of the Configurator.
- 4 Set Protocol=0
- 5 Clear the Match only established TCP connections check box.
- 6 Set Port=Ignore
- 7 Set Mask (Source)=255.255.255.192
- 8 Set Address (Source)=192.100.40.128
- 9 Set Port (Destination)=Ignore
- 10 Set Mask (Destination)=0.0.0.0
- 11 Set Address (Destination)=0.0.0.0
- 12 Set Filter Action to Forward.

Saving the filter

To actually create the filter you have defined, see “Saving the settings” on page 13-3.

A sample IP filter for more complex security issues

This section describes an IP filter that illustrates some of the issues you might need to consider when writing your own IP filters. The sample filter does not address intricate points of network security. You might want to use it 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 WAN access to the server’s IP address while restricting WAN access to 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 must be directed appropriately to the originating host. In this example, the Web server’s IP address is 192.9.250.5.

To define this IP filter, you first set it up, then define conditions to:

- Forward packets destined for the web server.
- Forward responses to TCP requests.
- Forward UDP packets.
- Forward IXMP packets.

Setting up the filter

To set this filter up:

- 1 On the left side of the Configurator, click the Filters button.
The Add/Copy/Delete window appears in the lower-left section of the Configurator.
- 2 On the Add/Copy/Delete menu, click Add.
A dialog box appears and prompts you for a value to apply as Filter Name.
- 3 Enter a name for the filter, and select OK. For example:

Filter Name=WebOK

- 4 Set Edit Condition to 1.
- 5 Click the Input Filters tab.
- 6 Select the Enabled check box on the right side of the Configurator.
- 7 Click the IP Packets button.

Forwarding packets destined for the web server

Configure the Input condition one to forward any packet with a web server's IP address as its destination:

- 1 Set Protocol=6
Protocol 6 specifies TCP packets.
- 2 Clear the "Match only established TCP connections" check box.
- 3 Set Port=Ignore
- 4 Set Mask (Source)=0.0.0.0
- 5 Set Address (Source)=0.0.0.0
- 6 Set Port (Destination)=Equal to
- 7 Set Port number to 80.
- 8 Set Mask (Destination)=255.255.255.255
The Pipeline 220 will compare all 32 bits of the destination IP address.
- 9 Set Address (Destination)=192.9.250.5
This is the IP address of the web server.
- 10 Set Filter Action to Forward.

Forwarding responses to TCP requests

Configure a Input condition two to forward any incoming TCP packet that uses a source port greater than 1023. The TCP protocol defines these packets as responses to TCP requests. The Pipeline 220 forwards these packets because the initial TCP request was generated by local devices. Local users are allowed to Telnet to remote devices, but remote devices are prevented from establishing TCP or Telnet connections to local devices. Proceed as follows:

- 1 Set Edit Condition to 2.
- 2 Select the Enabled check box on the right side of the Configurator.
- 3 Set Protocol=6
- 4 Clear the "Match only established TCP connections" check box.
- 5 Set Port (Source)=Ignore
- 6 Set Mask (Source)=0.0.0.0
- 7 Set Address (Source)=0.0.0.0
- 8 Set Port (Destination)=Greater than
- 9 Set Port number=1023
- 10 Set Mask (Destination)=0.0.0.0
- 11 Set Address (Destination)=0.0.0.0
- 12 Set Filter Action to Forward.

Forwarding UDP packets

Configure Input condition three to forward UDP packets. For example, a RIP packet is sent out as a UDP packet to destination port 520. The response to this RIP request is sent to a random destination port greater than 1023. Proceed as follows:

- 1 Set Edit Condition to 3.
- 2 Select the Enabled check box on the right side of the Configurator.
- 3 Set Protocol=17
Protocol 17 specifies UDP packets.
- 4 Clear the “Match only established TCP connections” check box.
- 5 Set Port (Source)=Ignore
- 6 Set Mask (Source)=0.0.0.0
- 7 Set Address (Source)=0.0.0.0
- 8 Set Port (Destination)=Greater than
- 9 Set Port number=1023
- 10 Set Mask (Destination)=0.0.0.0
- 11 Set Address (Destination)=0.0.0.0
- 12 Set Filter Action to Forward.

To complete the definition of this filter, configure Input condition four to forward ICMP packets, to allow unrestricted Pings and Traceroutes. ICMP does not use ports like TCP and UDP, so a source and destination port comparison is unnecessary. Proceed as follows:

- 1 Set Edit Condition to 4.
- 2 Select the Enabled check box on the right side of the Configurator.
- 3 Set Protocol=1
Protocol 1 specifies UDP packets.
- 4 Clear the “Match only established TCP connections” check box.
- 5 Set Port (Source)=Ignore
- 6 Set Mask (Source)=0.0.0.0
- 7 Set Address (Source)=0.0.0.0
- 8 Set Port (Destination)=Ignore
- 9 Set Mask (Destination)=0.0.0.0
- 10 Set Address (Destination)=0.0.0.0
- 11 Set Filter Action to Forward.

Saving the filter

To actually create the filter you have defined, see “Saving the settings” on page 13-3.

Setting Up Virtual Private Networking

14

This chapter covers the following topics:

Introduction to Virtual Private Networking (VPN)	14-1
Configuring ATMP tunnels	14-1

Introduction to Virtual Private Networking (VPN)

Virtual Private Networks provide low-cost remote access to private LANs via the Internet. The tunnel to the private corporate network can be from an ISP, enabling mobile nodes to dial-in to a corporate network, or it can provide a low-cost Internet connection between two corporate networks. Ascend currently supports two VPN schemes: Ascend Tunnel Management Protocol (ATMP) and Point-to-Point Tunneling Protocol (PPTP).

An ATMP session occurs between two Ascend units via UDP/IP. All packets passing through the tunnel are encapsulated in standard GRE (Generic Routing Encapsulation) as described in RFC 1701. ATMP creates and tears down a cross-Internet tunnel between the two Ascend units. In effect, the tunnel collapses the Internet cloud and provides what looks like direct access to a home network. Bridging is not supported through the tunnels. All packets must be routed with IP or IPX.

Point-to-Point-Tunneling Protocol (PPTP) was developed by Microsoft Corporation to enable Windows 95 and Windows NT Workstation users to dial into a local ISP to connect to a private corporate network across the Internet.

The Pipeline 220 does not support dial-in users, so its support of PPTP consists of routing or forwarding PPTP traffic as appropriate. The Pipeline 220 does not act as either a PPTP Access Concentrator (PAC) or a PPTP Network Server (PNS).

Configuring ATMP tunnels

This section describes how ATMP tunnels work between an Ascend MAX and a Pipeline 220. The MAX is configured as a *foreign* agent (typically a local ISP) and the Pipeline 220 as a *home* agent, with access to the home network. A mobile node dials into the foreign agent, which establishes a cross-Internet IP connection to the home agent. The foreign agent then requests an ATMP tunnel on top of the IP connection.

The home agent is the terminating part of the tunnel, where most of the ATMP processing occurs. It communicates with the home network (the destination network for mobile nodes) through a direct connection, another router, or across a nailed connection.

For example, in Figure 14-1, the mobile node might be a sales person who logs into an ISP to access his or her home network. The ISP is the foreign agent. The home agent has access to the home network.

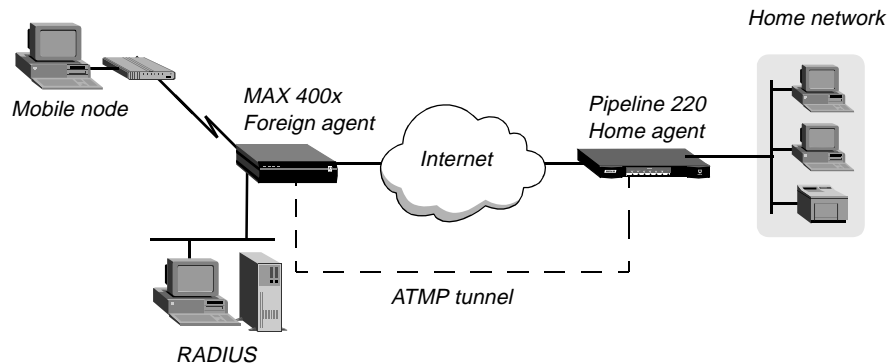


Figure 14-1. ATMP tunnel across the Internet

How the Pipeline 220 creates ATMP tunnels

The mobile node, foreign agent, and home agent establish an ATMP-tunnel connection as follows:

- 1 A mobile node dials a connection to the foreign agent.
- 2 After successful authentication, the foreign agent communicates with the home agent, and an IP connection establishes.
- 3 The foreign agent informs the home agent that the mobile node is connected, and requests a tunnel. It sends up to 10 RegisterRequest messages at 2-second intervals, timing out and logging a message if it receives no response to those requests.
- 4 The home agent requests a password before it creates the tunnel.
- 5 The foreign agent returns an encrypted version of the Ascend-Home-Agent-Password found in the mobile node's RADIUS profile. This password must match the home agent's Password parameter in the ATMP configuration in the Ethernet Profile.
- 6 The home agent returns a RegisterReply with a number that identifies the tunnel. If registration fails, the foreign agent disconnects the mobile node. If registration succeeds, the home agent creates the tunnel between itself and the foreign agent.
- 7 When the mobile node disconnects from the foreign agent, the foreign agent sends a DeregisterRequest to the home agent to close down the tunnel.
The foreign agent can send its request a maximum of ten times, or until it receives a DeregisterReply. If the foreign agent receives packets for a mobile node whose connection has been terminated, the foreign agent silently discards the packets.

Router and gateway mode

The home agent can communicate with the home network through a direct connection, through another router, or across a nailed connection. When the home agent relies on packet routing to reach the home network, it operates in router mode. When it has a nailed connection to the home network, it is in gateway mode.

Configuring a home agent in router mode

When the ATMP tunnel has been established between the home agent and foreign agent, the home agent in router mode 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.

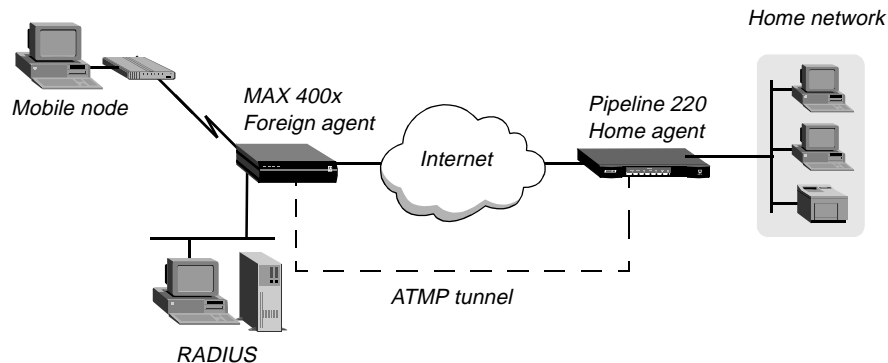


Figure 14-2. Home agent routing to the home network

Understanding the ATMP router mode parameters

This section provides some background information about parameters used in configuring the Pipeline 220 as a home agent in router mode:

Parameter	How it's used
Agent Mode	To enable the Pipeline 220's home-agent functionality, Agent Mode must be set to Home.
Home Agent Type	When Home Agent Type is set to Router, the home agent relies on routing (not a WAN connection) to pass packets, received through the tunnel, to the home network.
Home Agent Password	This is the password used to authenticate the ATMP tunnel itself. It must match the password specified in the Ascend-Home-Agent-Password attribute of mobile nodes' RADIUS profiles.
Local UDP port	By default, ATMP uses UDP port 5150 for ATMP messages between the foreign and home agents. If you specify a different UDP port number, make sure it is consistent between the home agent and foreign agent.
SAP Reply	Enables the home agent to reply to the mobile node's IPX Nearest Server Query if it knows about a server on the home network. If set to No, the home agent simply tunnels the mobile node's request to the home network.
IP configuration and Connection profile	The cross-Internet connection to the foreign agent is an IP routing connection, which is authenticated and established in the usual way. For details, see Chapter 6, "Configuring IP Routing."

Notes about routing to the mobile node

When the home agent receives IP packets through the ATMP tunnel, it adds a host route for the mobile node to its IP routing table. Then it supports IP and IPX routing normally. When the home agent receives IPX packets through the tunnel, it adds a route to the mobile node on the basis of the virtual IPX network number assigned in the RADIUS user profile.

For IP routes, you can enable RIP on the home agent's Ethernet to enable other hosts and networks to route to the mobile node. Enabling RIP is particularly useful if the home network is one or more hops away from the home agent's Ethernet. If RIP is turned off, other routers require static routes that specify the home agent as the route to the mobile node.

Note: If the home agent's Ethernet is the home network (a direct connection), you should turn on proxy ARP in the home agent so that local hosts can use ARP to find the mobile node.

For details about IP routes, see Chapter 6, "Configuring IP Routing." For information about IPX routes, see Chapter 10, "Configuring IPX Routing."

Example of configuring a home agent in router mode (IP)

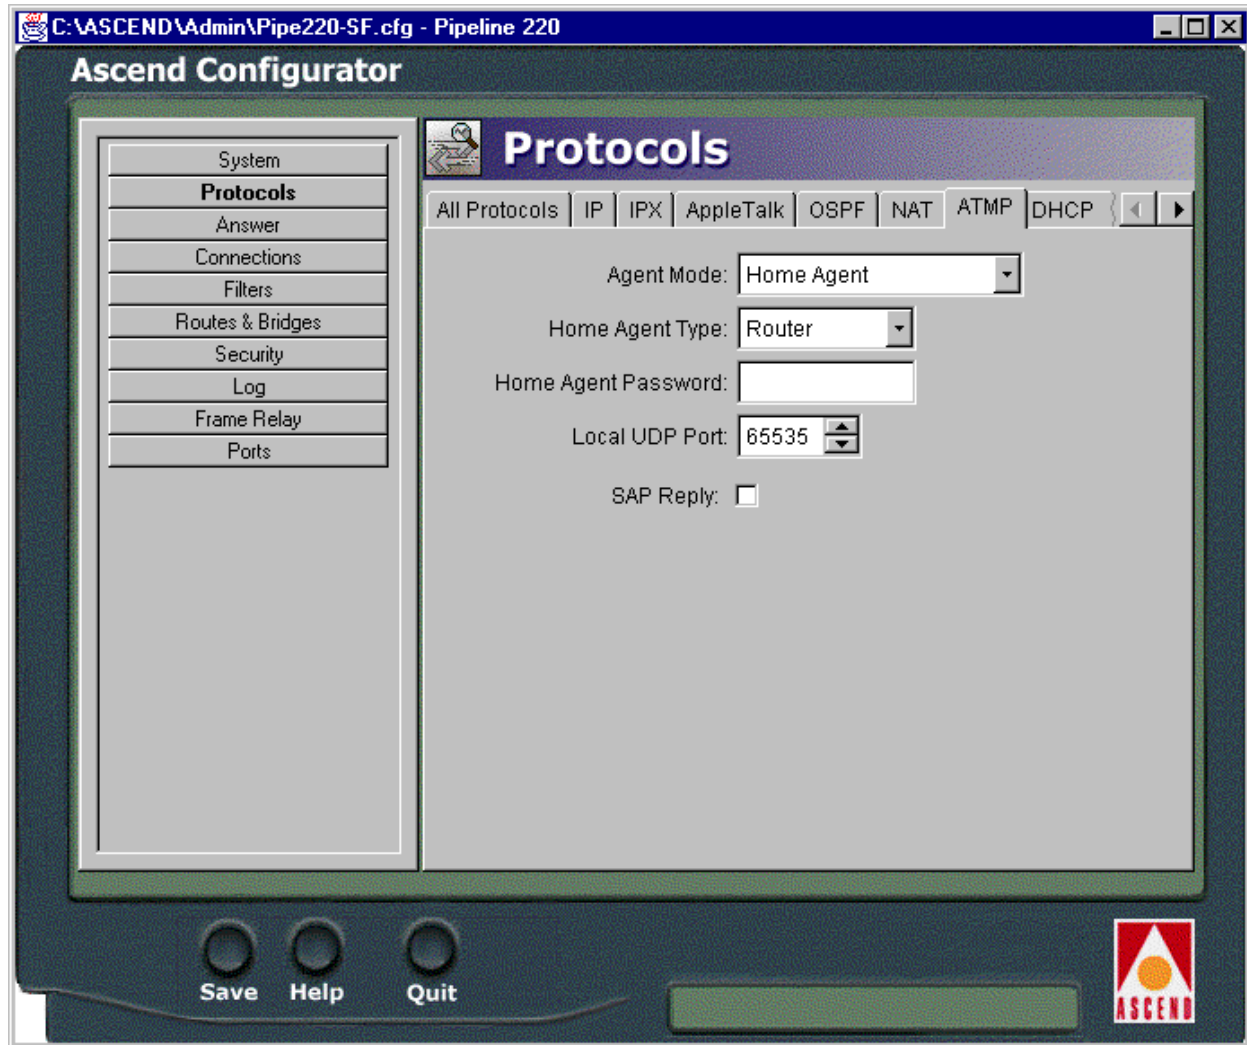
Before configuring it as a home agent in IP router mode, verify that the Pipeline 220 has a valid IP address in the Protocols button > IP tab > Addresses button > IP Address parameter. You should also validate IP connectivity, by pinging the Pipeline 220 from another IP host.

Configuring system-wide ATMP parameters

To begin configuring the home agent in router mode to reach an IP home network:

- 1 On the left side of the Configurator, click the Protocols button.

- 2 On the right side of the Configurator, click the ATMP tab.
System-wide ATMP parameters appear:



- 3 Set Agent Mode to Home.
- 4 Set Home Agent Type to Router.
- 5 Set Password to the value that will be supplied by the RADIUS profile of mobile users.
The attribute in the mobile user's profile is Ascend-Home-Agent-Password (unlike Dial-In Password, Ascend-Home-Agent-Password is the same for all users).

Configuring a Connection profile to the foreign agent


To configure a Connection profile to provide a route to the foreign agent:

- 1 On the left side of the Configurator, click the Connections button.
The Add/Copy/Delete window appears in the lower-left section of the Configurator.
- 2 In the Add/Copy/Delete window, click Add.

Setting Up Virtual Private Networking

Configuring ATMP tunnels

The Configurator displays a dialog box prompting you for the name of the new connection:



- 3 Enter a name and click the OK button.

The new profile appears on the left side of the Configurator, with default values for parameters:



- 4 Set the connection parameters as required for your connection. Be sure to include the following settings:
 - Select the Connection button > General tab > Active check box.
 - Select the Connection button > IP tab > Enable IP Routing check box.
 - Set the Connection button > IP tab > Addresses button > IP Address parameter to the foreign agent's Ethernet interface IP address.

- Set the Connection button > Authentication tab > Authentication parameter to the type of authentication used.
- Set the Connection button > Authentication tab > Dial-In Password parameter to the value specified in the Mobile user's profile.

Note: The Dial-In Password should be unique for every mobile user, whereas the Password for the ATMP tunnel (in the Protocols button > ATMP tab > Password parameter) must be identical for all mobile users, and specified in the Ascend-Home-Agent-Password attribute for each user.

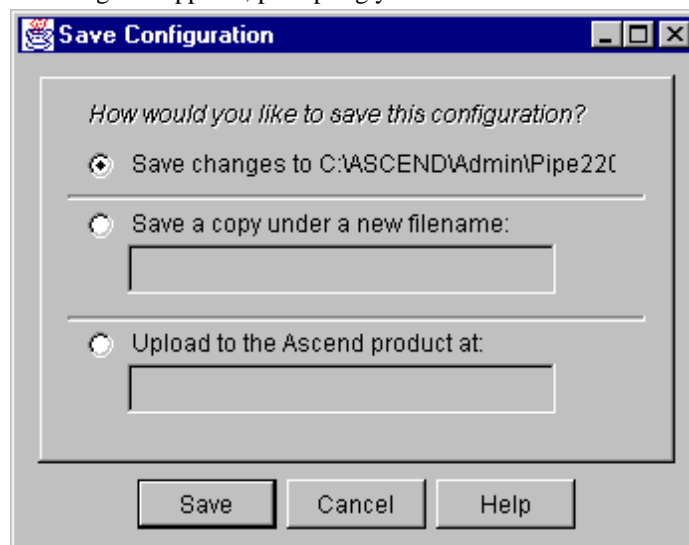
- 5 Save the new settings, as described in "Saving the settings".

Saving the settings

The Pipeline 220 settings you configure with the Ascend Configurator do not take effect until you upload them to the Pipeline 220. You use the Save command to upload the changes. Alternatively, you can save the setting to a text file for subsequent use, or to another device. To save the settings:

In the lower-left corner of the Configurator, click Save

- 1 A dialog box appears, prompting you for a save method:



- 2 Select one of the following:

- If you are uploading the configuration to the same device, select "Upload changes to <Pipeline 220 name>."
- If you are saving the configuration to a text file on your computer, select "Save a copy under a new filename" and enter the filename to which the configuration should be saved.
- If you are uploading the configuration to a different device, select "Upload to the Ascend product at:" and enter the IP address of the device.

- 3 Click Save.

If you selected "Upload changes to <Pipeline 220 name>" in step 2, the Ascend Configurator uploads the configuration file to the Pipeline 220 via TFTP. If you selected

one of the other options, the Configurator sends the file to the location you specified. If you specified a remote device, transmission is via TFTP.

Example of configuring a home agent in router mode (IPX)

Before configuring it as a home agent in IPX router mode, verify that the Pipeline 220 must be configured to route IPX. For details, see Chapter 10, “Configuring IPX Routing.”

Configuring system-wide ATMP parameters

To begin configuring the home agent in router mode to reach an IPX network:

- 1 On the left side of the Configurator, click the Protocols button.
- 2 On the right side of the Configurator, click the ATMP tab.
- 3 Set Agent Mode to Home.
- 4 Set Home Agent Type to Router.
- 5 Set Password to the value that will be supplied by the RADIUS profile of mobile users.
The attribute in the mobile user’s profile is Ascend-Home-Agent-Password (unlike Dial-In Password, Ascend-Home-Agent-Password is the same for all users).
- 6 Select the SAP Reply check box.

Configuring a Connection profile to the foreign agent

Now configure a Connection profile to provide a route to the foreign agent:

- 1 On the left side of the Configurator, click the Connections button.
The Add/Copy/Delete window appears in the lower-left section of the Configurator.
- 2 In the Add/Copy/Delete window, click Add.
The Configurator displays a dialog box prompting you for the name of the new connection.
- 3 Enter a name and click the OK button.
The new profile appears on the left side of the Configurator, with default values for parameters.
- 4 Set the connection parameters required for your connection. Be sure to include the following settings:
 - Select the Connection button > General tab > Active check box.
 - Select the Connection button > IP tab > Enable IP Routing check box.
 - Set the Connection button > IP tab > Addresses button > IP Address parameter to the foreign agent’s Ethernet interface IP address.
 - Set the Connection button > Authentication tab > Authentication parameter to the type of authentication used.
 - Set the Connection button > Authentication tab > Dial-In Password parameter to the value specified in the Mobile user’s profile.

Note: The Dial-In Password should be unique for every mobile user, whereas the Password for the ATMP tunnel (in the Protocols button > ATMP tab > Password parameter) must be identical for all mobile users, and specified in the Ascend-Home-Agent-Password attribute for each user.

- 5 Save the new settings, as described in “Saving the settings” .

Saving the settings

The Pipeline 220 settings you configure with the Ascend Configurator do not take effect until you upload them to the Pipeline 220. You use the Save command to upload the changes. Alternatively, you can save the setting to a text file for subsequent use, or to another device. To save the settings:

- 1 In the lower-left corner of the Configurator, click Save
- 2 A dialog box appears, prompting you for a save method. Select one of the following:
 - If you are uploading the configuration to the same device, select “Upload changes to <Pipeline 220 *name*>.”
 - If you are saving the configuration to a text file on your computer, select “Save a copy under a new filename” and enter the filename to which the configuration should be saved.
 - If you are uploading the configuration to a different device, select “Upload to the Ascend product at:” and enter the IP address of the device.

- 3 Click Save.

If you selected “Upload changes to <Pipeline 220 *name*>” in step 2, the Ascend Configurator uploads the configuration file to the Pipeline 220 via TFTP. If you selected one of the other options, the Configurator sends the file to the location you specified. If you specified a remote device, transmission is via TFTP.

Configuring a home agent in gateway mode

When the home agent is configured in gateway mode, it receives GRE-encapsulated IP packets from the foreign agent, strips off the encapsulation, and passes the packets across a nailed WAN connection to the home network.

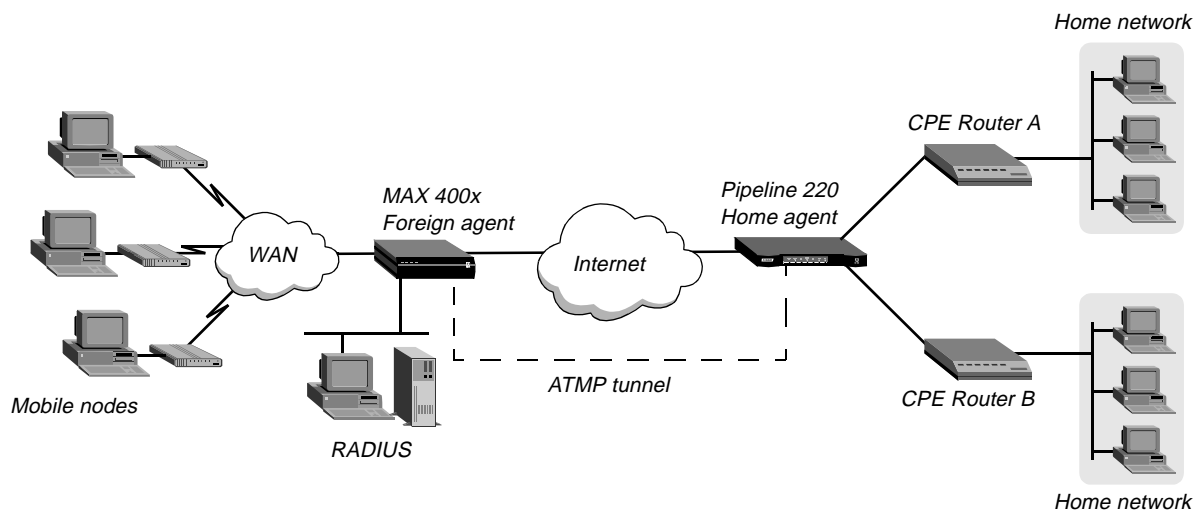


Figure 14-3. Home agent in gateway mode

Note: To enable hosts and routers on the home network to reach the mobile node, you must configure a static route in the Customer Premise Equipment (CPE) router on the home network (not in the home agent). The static route must specify the home agent as the route to the mobile node. That is, the route's destination address specifies the Framed-Address of the mobile node, and its gateway address specifies the IP address of the home agent.

Understanding the ATMP gateway mode parameters

This section provides some background information about parameters used in configuring the Pipeline 220 as a home agent in gateway mode:

Parameter	How it is used
Agent Mode	To enable the Pipeline 220's home agent functionality, Agent Mode must be set to Home.
Home Agent Type	When the Home Agent Type is set to Gateway, the home agent forwards packets received through the tunnel to the home network across a nailed WAN connection.
Home Agent Password	This is the password used to authenticate the ATMP tunnel itself. It must match the password specified in the Ascend-Home-Agent-Password attribute of mobile nodes' RADIUS profiles.
Local UDP port	By default, ATMP uses UDP port 5150 for ATMP messages between the foreign and home agents. If you specify a different UDP port number, make sure it is consistent between the home agent and foreign agent.
SAP Reply	Enables the home agent to reply to the mobile node's IPX Nearest Server Query if it knows about a server on the home network. If set to No, the home agent simply tunnels the mobile node's request to the home network.
IP configuration and Connection profile	The cross-Internet connection to the foreign agent is an IP routing connection, which is authenticated and established in the usual way. For details, see Chapter 6, "Configuring IP Routing."
Connection profile to the home network	The Connection profile to the home network must be a local profile. It cannot be specified in RADIUS. The name of this Connection profile must match the name in the Ascend-Home-Network-Name attribute in the mobile node's RADIUS profile.

Example of configuring a home agent in gateway mode (IP)

Before configuring it as a home agent in gateway mode, verify that the Pipeline 220 has a valid IP address in the Protocols button > IP tab > Addresses button > IP Address parameter. You should also validate IP connectivity, by pinging the Pipeline 220 from another IP host.

Configuring system-wide ATMP parameters

To configure the home agent in gateway mode to reach an IP home network:

- 1 On the left side of the Configurator, click the Protocols button.

- 2 On the right side of the Configurator, click the ATMP tab.
- 3 Set Agent Mode to Home.
- 4 Set Home Agent Type to Gateway.
- 5 Set Password to the value that will be supplied by the RADIUS profile of mobile users.
The attribute in the mobile user's profile is Ascend-Home-Agent-Password (unlike Dial-In Password, Ascend-Home-Agent-Password is the same for all users).

Configuring a Connection profile to the foreign agent

To configure a Connection profile to provide a route to the foreign agent:

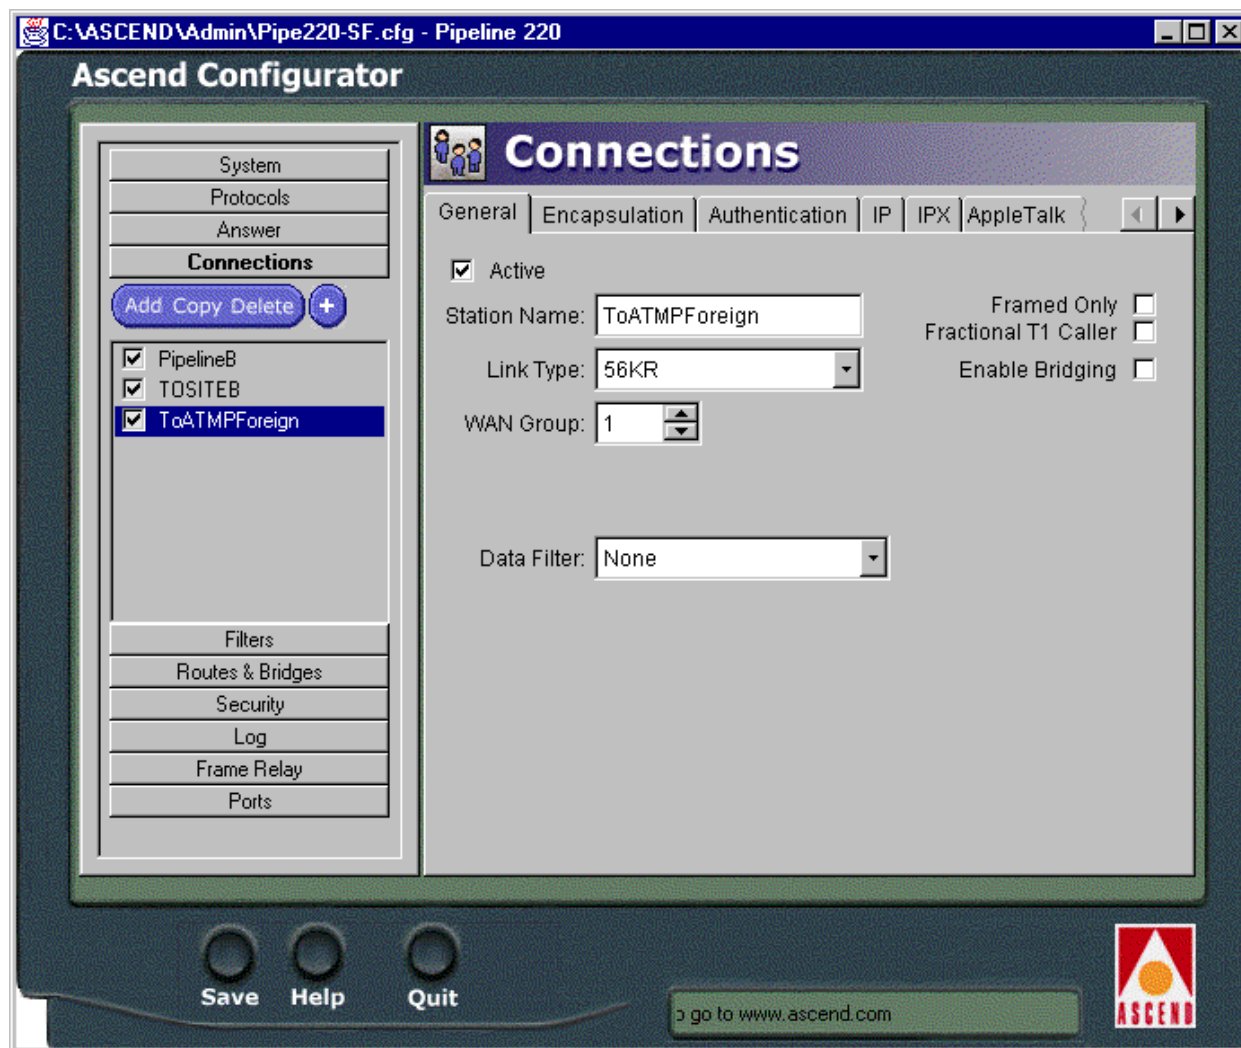
- 1 On the left side of the Configurator, click the Connections button.
The Add/Copy/Delete window appears in the lower-left section of the Configurator.
- 2 In the Add/Copy/Delete window, click Add.
The Configurator displays a dialog box prompting you for the name of the new connection.

Setting Up Virtual Private Networking

Configuring ATMP tunnels

- 3 Enter a name and click the OK button.

The new profile appears on the left side of the Configurator, with default values for parameters:



- 4 Set the connection parameters as required for your connection. Be sure to include the following settings:
 - Select the Connection button > General tab > Active check box.
 - Select the Connection button > IP tab > Enable IP Routing check box.
 - Set the Connection button > IP tab > Addresses button > IP Address parameter to the foreign agent's Ethernet interface IP address.
 - Set the Connection button > Authentication tab > Authentication parameter to the type of authentication used.
 - Set the Connection button > Authentication tab > Dial-In Password parameter to the value specified in the Mobile user's profile.

Configuring a Connection profile to the home network

To configure a connection profile for the nailed WAN link to the home network:

- 1 On the left side of the Configurator, click the Connections button.
The Add/Copy/Delete window appears in the lower-left section of the Configurator.
- 2 In the Add/Copy/Delete window, click Add.
The Configurator displays a dialog box prompting you for the name of the new connection.
- 3 Enter a name and click the OK button.
The new profile appears on the left side of the Configurator, with default values for parameters.
- 4 Set the connection parameters as required for your connection. Be sure to include the following settings:
 - Select the Connection button > General tab > Active check box.
 - Select the Connection button > General tab > ATMP Gateway parameter check box.
 - Select the Connection button > IP tab > Enable IP Routing check box.
 - Set the Connection button > IP tab > Addresses button > IP Address parameter to the IP address of the home network.
- 5 Save the new settings, as described in “Saving the settings”

Saving the settings

The Pipeline 220 settings you configure with the Ascend Configurator do not take effect until you upload them to the Pipeline 220. You use the Save command to upload the changes. Alternatively, you can save the setting to a text file for subsequent use, or to another device. To save the settings:

- 1 In the lower-left corner of the Configurator, click Save
- 2 A dialog box appears, prompting you for a save method. Select one of the following:
 - If you are uploading the configuration to the same device, select “Upload changes to <Pipeline 220 name>.”
 - If you are saving the configuration to a text file on your computer, select “Save a copy under a new filename” and enter the filename to which the configuration should be saved.
 - If you are uploading the configuration to a different device, select “Upload to the Ascend product at:” and enter the IP address of the device.
- 3 Click Save.
If you selected “Upload changes to <Pipeline 220 name>” in step 2, the Ascend Configurator uploads the configuration file to the Pipeline 220 via TFTP. If you selected one of the other options, the Configurator sends the file to the location you specified. If you specified a remote device, transmission is via TFTP.

Example of configuring a home agent in gateway mode (IPX)

Before configuring it as a home agent in IPX router mode, verify that the Pipeline 220 is configured to route IPX. For additional details, see Chapter 10, “Configuring IPX Routing.”

Configuring system-wide ATMP parameters

To begin configuring the home agent in gateway mode to reach an IPX network:

- 1 On the left side of the Configurator, click the Protocols button.
- 2 On the right side of the Configurator, click the ATMP tab.
- 3 Set Agent Mode to Home.
- 4 Set Home Agent Type to Gateway.
- 5 Set Password to the value that will be supplied by the RADIUS profile of mobile users.
The attribute in the mobile user's profile is Ascend-Home-Agent-Password (unlike Dial-In Password, Ascend-Home-Agent-Password is the same for all users).
- 6 Select the SAP Reply check box.

Configuring a Connection profile to the foreign agent

To configure a Connection profile to provide a route to the foreign agent:

- 1 On the left side of the Configurator, click the Connections button.
The Add/Copy/Delete window appears in the lower-left section of the Configurator.
- 2 In the Add/Copy/Delete window, click Add.
The Configurator displays a dialog box prompting you for the name of the new connection.
- 3 Enter a name and click the OK button.
The new profile appears on the left side of the Configurator, with default values for parameters.
- 4 Set the connection parameters as required for your connection. Be sure to include the following settings:
 - Select the Connection button > General tab > Active check box.
 - Select the Connection button > IP tab > Enable IP Routing check box.
 - Set the Connection button > IP tab > Addresses button > IP Address parameter to the foreign agent's Ethernet interface IP address.
 - Set the Connection button > Authentication tab > Authentication parameter to the type of authentication used.
 - Set the Connection button > Authentication tab > Dial-In Password parameter to the value specified in the Mobile user's profile.

Configuring a Connection profile to the home network

To configure a connection profile for the nailed WAN link to the home network:

- 1 On the left side of the Configurator, click the Connections button.
The Add/Copy/Delete window appears in the lower-left section of the Configurator.
- 2 In the Add/Copy/Delete window, click Add.
The Configurator displays a dialog box prompting you for the name of the new connection.
- 3 Enter a name and click the OK button.

The new profile appears on the left side of the Configurator, with default values for parameters.

- 4** Set the connection parameters as required for your connection. Be sure to include the following settings:
 - Select the Connection button > General tab > Active check box.
 - Select the Connection button > General tab > ATMP Gateway parameter check box.
 - Select the Connection button > IP tab > Enable IP Routing check box.
 - Select the Connection button > IPX tab > Enable IPX Routing check box.
 - Set the Connection button > IPX tab > RIP Updates parameter to Send & Receive.
 - Set the Connection button > IPX tab > SAP parameter to Send & Receive.
 - Set the Connection button > IP tab > Addresses button > IP Address parameter to the IP address of the home network.
- 5** Save the new settings, as described in “Saving the settings”.

Saving the settings

The Pipeline 220 settings you configure with the Ascend Configurator do not take effect until you upload them to the Pipeline 220. You use the Save command to upload the changes. Alternatively, you can save the setting to a text file for subsequent use, or to another device. To save the settings:

- 1** In the lower-left corner of the Configurator, click Save
- 2** A dialog box appears, prompting you for a save method. Select one of the following:
 - If you are uploading the configuration to the same device, select “Upload changes to <Pipeline 220 name>.”
 - If you are saving the configuration to a text file on your computer, select “Save a copy under a new filename” and enter the filename to which the configuration should be saved.
 - If you are uploading the configuration to a different device, select “Upload to the Ascend product at:” and enter the IP address of the device.
- 3** Click Save.

If you selected “Upload changes to <Pipeline 220 name>” in step 2, the Ascend Configurator uploads the configuration file to the Pipeline 220 via TFTP. If you selected one of the other options, the Configurator sends the file to the location you specified. If you specified a remote device, transmission is via TFTP.

SNMP administrative support

15

This chapter covers the following topics:

Introduction	15-1
Configuring SNMP access security	15-1
Setting SNMP traps	15-5
Ascend Enterprise traps	15-8
Supported MIBs	15-10

Introduction

The Pipeline 220 supports SNMP on an IP-routed network. An SNMP management station that uses the Ascend Enterprise Management Information Base (MIB) can request information from the Pipeline 220, set parameters, and send alarm notifications when specific conditions occur in the Pipeline 220. An SNMP manager must be running on a host on the local IP network, and the Pipeline 220 must be able to find that host, either via a static route or RIP updates.

SNMP supports password security, which you should configure to protect the Pipeline 220 from modification by unauthorized users with access to SNMP management stations.

SNMP traps provide SNMP management stations with real-time system changes. Traps are messages sent notifying SNMP managers of specific events. For example, the Pipeline 220 can notify the SNMP management station that the condition of its nailed link has changed.

Configuring SNMP access security

There are two levels of SNMP security:

- Community strings, which are passwords that you must know for access to the Pipeline 220.
- Address security, which excludes SNMP access unless it is initiated from a specified IP address.

How the SNMP security options work

The Configurator's System button provides access to the parameters for configuring community strings and address security.

Community strings

The System button > SNMP tab > “Read Community Name” parameter specifies the SNMP community name allowing read access only. The System button > SNMP tab > “Read-Write Community” parameter specifies the SNMP community name for read and write access.

Address security

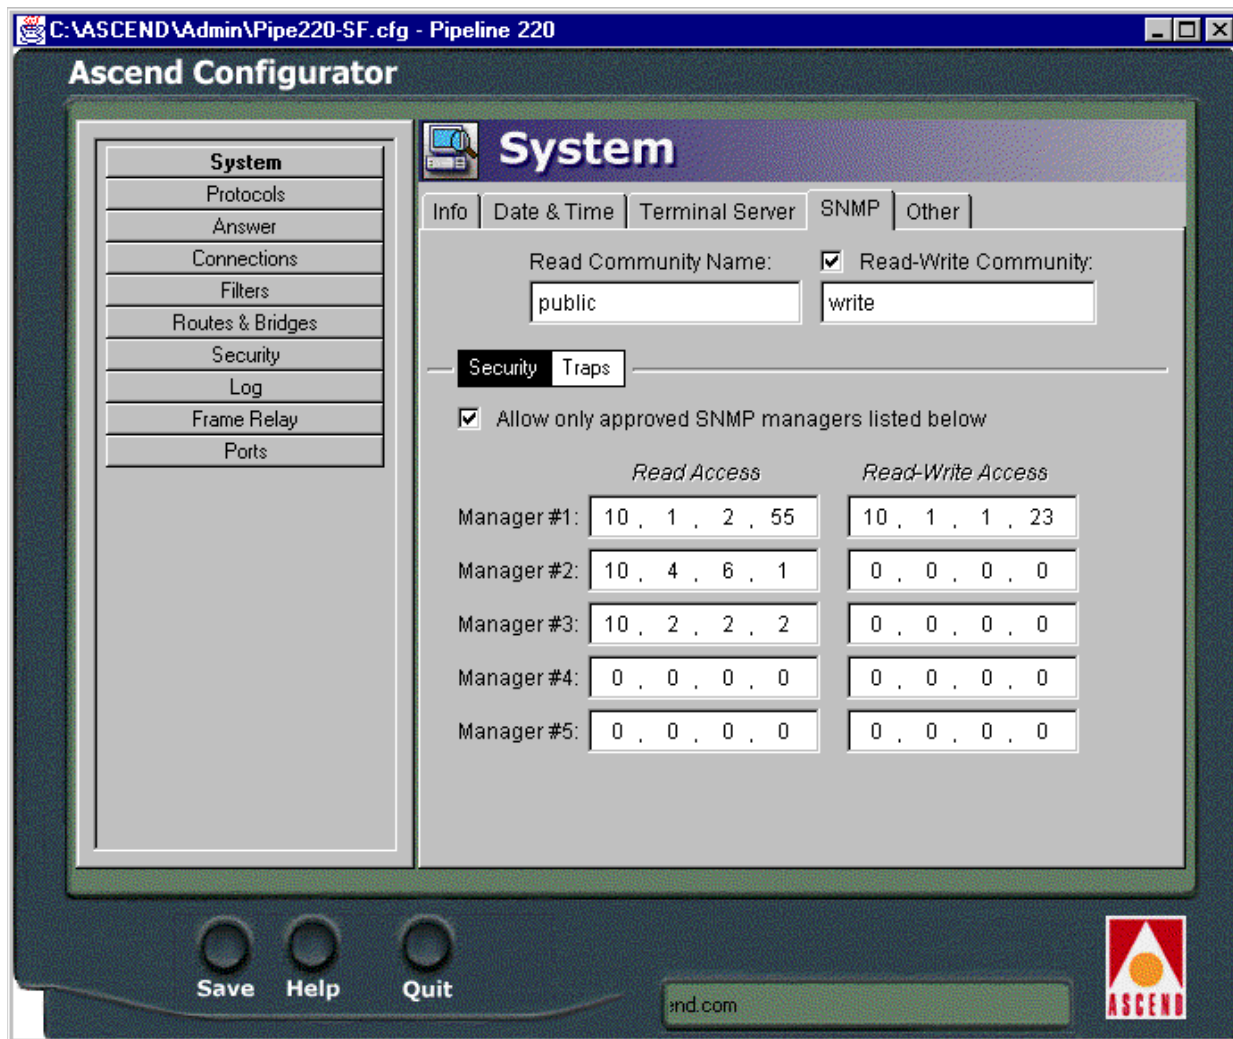
If you clear the System button > SNMP tab > Security button > “Allow only approved SNMP managers listed below” check box (the default value), any SNMP manager that supplies the configured community name is allowed access to the Pipeline 220. If you select the check box, the Pipeline 220 allows access only to SNMP managers whose IP addresses are listed in the System button > SNMP tab > Security button > Security button > “Manager #” parameters. You can specify up to five addresses.

Entering SNMP security settings

You can carry out the following procedure to set the community strings, enforces address security, and prevent write access:

- 1** On the left side of the Configurator, click the System button.
- 2** On the right side of the Configurator, click the SNMP tab.
Security and Traps buttons appear on the right side of the Configurator.


- 3 On the right side of the Configurator, click the Security button.
Security options appear on the right side of the Configurator:



- 4 Specify the Read Community Name and Read/Write Community parameter strings.
- 5 Select the “Allow only approved SNMP managers listed below” check box.
- 6 Specify up to five host addresses allowing Read access, and up to five host addresses allowing both Read and Write access.

You have now entered the security settings, but they do not take effect until you upload them to the Pipeline 220. You use the Save command to upload the changes. Alternatively, you can save the setting to a text file for subsequent use, or to another device.

- 7 In the lower-left corner of the Configurator, click Save
A dialog box appears, prompting you for a save method:



- 8 Select one of the options, as follows:
- If you are uploading the configuration to the same device, select "Upload changes to <Pipeline 220 name>."
 - If you are saving the configuration to a text file on your computer, select "Save a copy under a new filename" and enter the filename to which the configuration should be saved.
 - If you are uploading the configuration to a different device, select "Upload to the Ascend product at:" and enter the IP address of the device.
- 9 Click Save.

If you selected "Upload changes to <Pipeline 220 name>" in step 8, the Ascend Configurator uploads the configuration file to the Pipeline 220 via TFTP. If you selected one of the other options, the Configurator sends the file to the location you specified. If you specified a remote device, transmission is via TFTP.

Setting SNMP traps

A trap is a mechanism for reporting system change in real time (for example, reporting a call coming into to a serial host port). When a trap is generated by some condition, a traps-PDU (protocol data unit) is sent across the Ethernet to the SNMP manager.

You can configure the Pipeline 220 with eight separate Trap profiles, directing it to send different combinations of traps to different SNMP managers. For redundancy, you can configure the Pipeline 220 to send identical combinations of traps to different SNMP managers.

Understanding the SNMP trap parameters

To set traps, you use the following parameters:

Parameter	Description
Name	Specifies the name of the Trap profile. If you configure more than one profile, attaching a descriptive name to the trap might help you organize and keep track of different profiles.
Community string	Used in communicating with the SNMP manager. It must contain the community name associated with the SNMP PDU.
Address of SNMP manager	Specifies the IP address of the system running the SNMP manager. Note: To prevent the Pipeline 220 from sending SNMP traps, set Dest=0.0.0.0.
Classes of traps to be sent to the specified host	Specify whether the Pipeline 220 traps alarm events, security events, and port events, and sends a trap-PDU to the SNMP manager.

Entering an SNMP trap configuration

You can configure up to eight separate Traps profiles. To configure an SNMP Trap profile:

- 1 On the left side of the Configurator, click the Systems button.
- 2 On the right side of the Configurator, click the SNMP tab.
Security and Traps buttons appear on the right side of the Configurator.

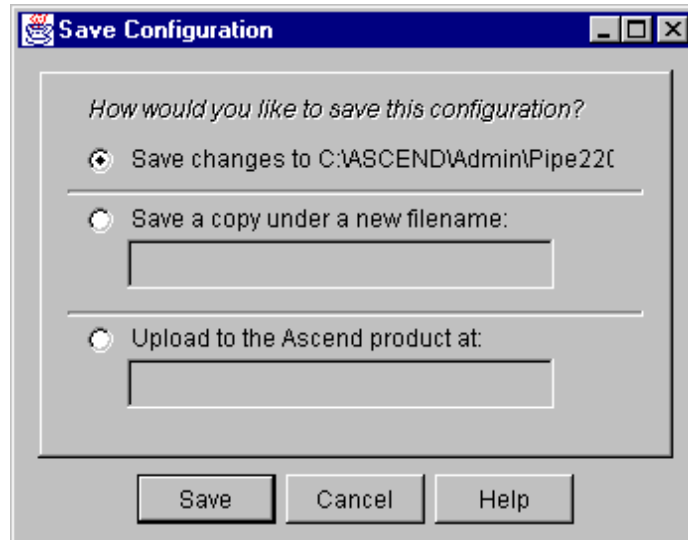
- 3 On the right side of the Configurator, click the Traps button.

Trap options appear in the right side of the Configurator:



- 4 Click the Traps: 1 button to configure the first Traps profile.
If this is not the first one, click the next available Traps button (up to eight).
- 5 Set the Name parameter with a descriptive name for the profile.
- 6 Specify the address of the SNMP manager to which the Pipeline 220 should send SNMP traps.
- 7 Set Community Name to the password used in communication with the SNMP manager.
The Pipeline 220 sends the value you specify to authenticate itself to the SNMP host when an SNMP trap event occurs.
- 8 Select the classes of traps to be sent to the specified SNMP manager.
You have now entered the settings for your Trap profile, but they do not take effect until you upload them to the Pipeline 220. You use the Save command to upload the changes. Alternatively, you can save the setting to a text file for subsequent use, or to another device.

- 9 In the lower-left corner of the Configurator, click Save
A dialog box appears, prompting you for a save method:



- 10 Select one of the following:
- If you are uploading the configuration to the same device, select "Upload changes to <Pipeline 220 name>."
 - If you are saving the configuration to a text file on your computer, select "Save a copy under a new filename" and enter the filename to which the configuration should be saved.
 - If you are uploading the configuration to a different device, select "Upload to the Ascend product at:" and enter the IP address of the device.
- 11 Click Save.

If you selected "Upload changes to <Pipeline 220 name>" in step 10, the Ascend Configurator uploads the configuration file to the Pipeline 220 via TFTP. If you selected one of the other options, the Configurator sends the file to the location you specified. If you specified a remote device, transmission is via TFTP.

Ascend Enterprise traps

This section is a brief summary of the traps generated by alarm, port, and security events. For detailed information, see the Ascend Enterprise MIB. For information about obtaining the Ascend MIB, see “Supported MIBs” on page 15-10.

Alarm events

Alarm events (also called *error events*) use trap types defined in RFC 1215 and 1315, in addition to an Ascend enterprise trap type. The Pipeline 220 supports the following trap types from RFC 1215 are supported:

Trap type	Description
coldStart (RFC-1215 trap-type 0)	Signifies that the Pipeline 220 sending the trap is reinitializing itself, so the configuration of the SNMP manager or the unit might be altered.
warmStart (RFC-1215 trap-type 1)	Signifies that the Pipeline 220 sending the trap is reinitializing itself, so that neither the configuration of SNMP manager or the unit is altered.
linkDown (RFC-1215 trap-type 2)	Signifies that the Pipeline 220 sending the trap recognizes a failure in one of the communication links represented in the SNMP manager's configuration.
linkUp (RFC-1215 trap-type 3)	Signifies that the Pipeline 220 sending the trap recognizes that one of the communication links represented in the SNMP manager's configuration has come up.
frDLCIStatusChange (RFC-1315 trap-type 1)	Signifies that the Pipeline 220 sending the trap recognizes that one of the virtual circuits (to which a DLCI number has been assigned) has changed state. That is, the link has either been created, invalidated, or it has toggled between the active and inactive states.
eventTableOverwrite (ascend trap-type 16)	Signifies that a new event has overwritten an unread event. This trap is sent only for systems that support Ascend's accounting MIB. Once sent, additional overwrites will not cause another trap to be sent until at least one table's worth of new events have occurred.

Port state change events

The following traps are effective on a port-by-port basis for each port pointed to by ifIndex. The hostPort objects are used to associate a change with ifIndex objects.

Trap type	Description
portInactive (ascend trap-type 0)	AIM port associated with the passed index has become inactive.
portDualDelay (ascend trap-type 1)	AIM port associated with the passed index is delaying the dialing of a second to avoid overloading devices that cannot handle two calls in close succession.

Trap type	Description
portWaitSerial (ascend trap-type 2)	AIM port associated with the passed index has detected DTR and is waiting for an HDLC controller to come online. CTS is off (V.25 bis dialing only).
portHaveSerial (ascend trap-type 3)	AIM port associated with the passed index is waiting for V.25 bis commands. CTS is on.
portRinging (ascend trap-type 4)	AIM port associated with the passed index has been notified of an incoming call.
portCollectDigits (ascend trap-type 5)	AIM port associated with the passed index is receiving digits from an RS366 interface (RS-366 dialing only).
portWaiting (ascend trap-type 6)	AIM port associated with the passed index is waiting for connect notification from the WAN after dialing or answer notification has been issued.
portConnected (ascend trap-type 7)	AIM port associated with the passed index has changed state. This change of state can be from connected to unconnected or vice versa. If connected to the far end, end-to-end data can flow but has not yet been enabled. The following trap report sequence shows a link is up: portWaiting (6) portConnected (7) portCarrier (8) The following trap report sequence shows a link is down: portConnected (7) portInactive (0)
portCarrier (ascend trap-type 8)	AIM port associated with the passed index has end-to-end data flow enabled.
portLoopback (ascend trap-type 9)	AIM port associated with the passed index has been placed in local loopback mode.
portAcrPending (ascend trap-type 10)	AIM port associated with the passed index has set ACR on the RS366 interface, and is waiting for the host device (RS-366 dialing only).
portDTENotReady (ascend trap-type 11)	AIM port associated with the passed index is waiting for DTE to signal a ready condition when performing X.21 dialing.

Security events

Security events are used to notify users of security problems, and to track access to the unit from the console. The MIB-II event *authenticationError* is a security event. The other security events are Ascend-specific.

Trap type	Description
authenticationFailure (RFC-1215 trap-type 4)	Signifies that the Pipeline 220 sending the trap is the addressee of a protocol message that is not properly authenticated.
consoleStateChange (ascend trap-type 12)	Signifies the console associated with the passed console index has changed state. To read the console's state, get ConsoleEntry from the Ascend enterprise MIB.

Trap type	Description
portUseExceeded (ascend trap-type 13)	The serial host port's use exceeds maximum set by Max DS0 Mins Port parameter associated with the passed index (namely, the interface number).
systemUseExceeded (ascend trap-type 14)	The serial host port's use exceeds maximum set by Max DS0 Mins System parameter associated with the passed index (namely, the interface number).
maxTelnetAttempts (ascend trap-type 15)	There have been three consecutive failed attempts to login onto the Pipeline 220 via Telnet.

Supported MIBs

You can download the most up-to-date version of the Ascend Enterprise MIB by logging in as *anonymous* to ftp.ascend.com. (No password is required.) In addition to the Ascend MIB, the Pipeline 220 also supports objects related to Ascend functionality in the following Internet standard MIBs:

- MIB-II implementation (RFC 1213)
- DS1 MIB implementation (RFC 1406)
- RS232 MIB implementation (RFC 1317)
- Frame Relay MIB implementation (RFC 1315)
- Modem MIB implementation (RFC 1696)

You can download the most recent version of the previous RFCs by logging in as *anonymous* to ftp.ds.internic.net. (No password is required.)

VT100 Interface System Administration

16

This chapter covers the following topics:

Introduction to Pipeline 220 administration.	16-1
Accessing the VT100 interface	16-2
Using the VT100 interface.	16-3
About Pipeline 220 passwords.	16-7
Using the Pipeline 220 status windows	16-8
Terminal-server command-line interface.	16-16

Introduction to Pipeline 220 administration

The Pipeline 220's VT100 interface provides a wide variety of features for monitoring and administering the unit's activities.

The initial display of the VT100 interface shows the Main Edit menu and a group of status windows. The status windows display a variety of information about the operation of your Pipeline 220. The Main Edit menu largely duplicates the functions of the Configurator, but does contain a few administrative commands not available through the Configurator. You also have access to DO commands, which enable you to perform additional tasks. (To perform any of the administrative tasks, you must activate administrative permissions.)

An additional advantage of being able to use the VT100 interface is that it provides access to the terminal-server command-line interface, which features a large assortment of powerful commands. For example: You can view the Pipeline 220 unit's routing tables and statistical information. You can access detailed information about the unit's IP routing table, OSPF routing table, and Frame Relay connections. You can also use the administrative commands Ping, Traceroute, Telnet, and IPXping to establish and test connectivity. You can manually add, delete or change routes in your IP routing table. Descriptions of the commands available through the terminal-server command-line interface form the major part of the this chapter.

Accessing the VT100 interface

You can access the VT100 user interface either through the Pipeline 220 Control port or via a Telnet connection over Ethernet. This section describes both access methods.

Using the Pipeline 220 control port

Before you can use the control port, some settings in your PC's communications software must match those on the Pipeline 220. Set the terminal emulation software as follows:

- 9600 bps
- 8 data bits
- No parity
- 1 stop bit
- No flow control
- Direct connect

To access the Pipeline 220's VT100 user interface via its control port:

- 1** Connect your PC's serial port to the Pipeline 220's control port with a serial cable.
- 2** On your PC, launch your communications software in terminal-emulation mode.
- 3** Press Ctrl-L to refresh the screen display.

The VT100 interface appears on your computer screen. The VT100 interface consists of a Main Edit menu on the left side of the display and eight status windows on the right side of the display.

Using Telnet

For Telnet access, the Pipeline 220 must have a valid IP address and be reachable by your PC over Ethernet. If you are unable to access the Pipeline 220 via Telnet, but believe it and the PC are configured correctly, contact your network administrator for help. To connect to the Pipeline 220's VT100 interface through Telnet:

- 1** From your PC, launch your Telnet software.
- 2** Telnet to the Pipeline 220's IP address. If prompted, enter the Telnet password.

The VT100 interface appears on your computer screen. The VT100 interface consists of a Main Edit menu on the left side of the display and eight status windows on the right side of the display.

Using the VT100 interface

This section explains how to navigate the VT100 interface.

Activating a menu or status window

You can interact with only one display at a time. The active display has a thick double line border on the left, right, and top sides.

If you press the Tab key, the thick double lines move to 00-200, the next screen to the right. If you continue pressing the Tab key, you activate each window from left to right and down, until you reach the last display in the lower right-hand corner. Back-Tab or Ctrl-O moves you in the opposite direction.

Opening menus and profiles

The Main Edit Menu contains a list of menus, each of which can contain profiles and submenus. In the menu that is currently open, the cursor character (>) points to one item in the menu. To move the cursor down, press Ctrl-N (next) or the down-arrow key. To move it up, press Ctrl-P (previous) or the up-arrow key. (Some VT100 emulators do not support the use of arrow keys.) For a complete list of key combinations used to navigate the interface, see Table 16-1 on page 16-6.

```
Main Edit Menu
  00-000 System
  >10-000 Serial Port T1-CSU
  20-000 Ethernet
```

To open a menu, move the cursor to the menu's name and press Enter. For example, press Ctrl-N until the cursor points to 30-000 Ethernet, and press Enter. The Ethernet menu opens.

```
30-000 Ethernet
  90-100 Connections
  90-200 Bridge Adrs
  90-300 Static Rtes
  90-400 Filters
  90-500 Firewalls
  90-600 Frame Relay
  90-700 Answer
  90-800 SNMP Traps
  90-900 IPX Routes
  90-A00 IPX SAP Filters
  90-B00 NAT
  90-C00 Mod Config
```

The Ethernet menu contains submenus and profiles related to network functionality, such as bridging, routing, WAN connections, and so forth. The Mod Config Profile in this menu relates to the configuration of the Ethernet interface itself, as shown next.

```
90-B00 Mod Config
  Module Name=
  Ether1 options...
```

```
Ether2 options...
WAN options...
SNMP options...
OSPF options...
OSPF global options...
Route Pref...
TServ options...
Bridging=No
IX Routing=No
Appletalk=No
Shared Prof=No
Telnet PW=****
RIP Policy=Poison Rvrs
RIP Summary=Yes
ICMP Redirects=Accept
```

Note: With the exception of parameters designated N/A (not applicable), you can edit all parameters in any profile. A profile is a group of parameters listed under a particular menu entry. N/A that means a parameter does not apply within the context of how some other parameter(s) or profile has been set.

Opening edit fields

To open an edit field for a text-based parameter (such as a password, for example), move the cursor to that parameter and press Enter. An edit field opens, delimited by brackets, as shown for the Telnet PW parameter, next.

```
90-B00 Mod Config
Module Name=
Ether1 options...
Ether2 options...
WAN options...
SNMP options...
OSPF options...
OSPF global options...
Route Pref...
TServ options...
Bridging=No
IX Routing=No
Appletalk=No
Shared Prof=No
Telnet PW:
[ ]

ICMP Redirects=Accept
```

Note: See “About Pipeline 220 passwords” on page 16-7 for related information.

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 and then delete or overwrite the characters.

To close the edit field and accept the new text, press Enter.

Setting enumerated parameters

An enumerated parameter is one for which there is a set of predefined values. You modify it by simply placing the cursor beside the parameter and typing the Enter, Return, or the Right-Arrow key until the proper value appears.

Saving your changes

When you exit a profile, you are prompted to confirm that you want to save changes.

```
EXIT?  
>0=ESC (Don't exit)  
1=Exit and discard  
2=Exit and accept
```

You can save the profile values by choosing the Exit and Save option and pressing Enter, or by pressing 2.

Special display characters and keys

The following characters have special meaning within the displays:

- The plus character (+) indicates that an input entry is too long to fit onto one line, and that the Pipeline 220 is truncating it for display purposes.
- Ellipses (...) mean that a submenu displays the details of a menu option.
The Pipeline 220 displays the submenu when you select the menu option.

The following table lists the special-purpose keys and key combinations you can use in the Control Monitor displays.

Table 16-1.Special purpose keys for Control Monitor displays (continued)

Table 16-2.

Key combination	Operation
Right-Arrow, Return, Enter, Ctrl-Z, Ctrl-F	Enumerated parameter: Select the next value. String value: Move one character to the right or enter the current input. Menu: Open the current selection.
Left-Arrow, Ctrl-X, Ctrl-B	Enumerated parameter: Select the previous value. String value: Move left one character or exit the current input. Menu: Close the current selection.
Down-Arrow, Ctrl-N	Move down to the next selection.
Up-Arrow, Ctrl-U, Ctrl-P	Move up to the previous selection.
Ctrl-V	Move to the next page of the list.
Tab, Ctrl-I	Move to the next window.
Back-Tab, Ctrl-O	Move to the previous window.
N/A	Toggle to a status menu from the edit menu and vice versa.
Delete	Delete the character under the cursor.
Backspace	Delete the character to the left of the cursor.
none	Overwrite the character under the cursor with a space.
Ctrl-D	Open the DO menu.
Ctrl-T	Return from or go to the Simplified Menus.
Ctrl-L	Refresh the VT-100 screen.
Ctrl-C	Return from the MIF to the normal menus.
D	Dial the currently selected profile.

Note: You always use the Control and Shift keys in combination with other keys. This document represents key combinations as two characters separated by a hyphen, such as Shift-T, which types the capital letter T.

About Pipeline 220 passwords

The Pipeline 220 has up to nine security levels, each of which is defined in a Security Profile. When shipped from the factory, all nine levels are wide open, with no defined restrictions. To see the list of Security Profiles, open the System menu in the Main Edit Menu, and then select Security and press Enter.

```
00-300 Security
>00-301 Default
    00-302
    00-303
    00-304
    00-305
    00-306
    00-307
    00-308
    00-309 Full Access
```

Whenever the Pipeline 220 is powered on, it activates the first Security Profile in this list, which is always named Default and always has no password.

Before you can use the administrative commands and profiles, you must log in as the superuser by activating a Security profile, such as the Full Access profile, that has sufficient permissions.

Note: For a session established via Telnet, you must first supply the Telnet password to establish a Telnet session. Then, the Default security level is set for that session. To configure the Pipeline 220 via Telnet, the user must activate the appropriate Security Profile.

To log in as the superuser, proceed as follows:

- 1 Press Ctrl-D to open the DO menu, then press P (or select P=Password).
- 2 In the list of Security profiles that opens, select Full Access.

The Pipeline 220 prompts you for the Full Access password. For example:

```
00-300 Security
Enter Password:
[ ]
```

Press > to accept

- 3 Type the password assigned to the profile and press Enter.

When you enter the correct password, the Pipeline 220 displays a message informing you that the password was accepted and that the Pipeline 220 is using the new security level:

```
Message #119
Password accepted.
Using new security level.
```

If the password you enter is incorrect, the Pipeline 220 prompts you again for the password.

Note: The default password for the Full Access login is *Ascend*.

One of the first thing most administrators do is to reset the privileges in the Default profile to restrict what can be done by anyone accessing the Pipeline 220 configuration menus. To do this:

- 1 Open the Default Security Profile and set the Operations privilege to No.
- 2 Assign a password to the Full Access Security Profile. (Do not restrict privileges in the Full Access Profile.)
- 3 Activate the Full Access Security Profile and proceed to configure the Pipeline 220.

Using the Pipeline 220 status windows

In the Pipeline 220 VT100 interface, the right side of the screen displays eight status windows. The status windows provide a great deal of read-only information about what is currently happening in the Pipeline 220.

The following section provides an overview of the information contained in the eight windows.

-----	-----
10-100 1234567890	20-200 Routes
L1/LA nnnnnnnnnn	>D: Default
12345678901234	G: 10.10.10.10
nnnnnnnnnnnnnn	LAN Active
-----	-----
20-100 Sessions	00-200 15:10:34
> 2 Active	>M31 Line Ch
0 davel-gw	LAN session up
0 lnemo.Ascend.COM	davel-gw
-----	-----
20-300 WAN Stat	90-400 Ether Stat
>Rx Pkt: 184318^	>Rx Pkt: 3486092
Tx Pkt: 159232	Tx Pkt: 10056
CRC: 0v	Col: 3530
-----	-----
00-100 Sys Option	20-700 Ether Opt
>Security Prof: 1 ^	>Enet I/F: UTP
Software +5.1Ap4+	Adr0: 00c07b6fd5b8
S/N: 5210003 v	Adr1: 00c07b6eadc0
-----	-----

Some of the status windows contain more information than can be displayed in the small window. If a lowercase v appears in the lower-right corner of a window, more information is available. To scroll through additional information in a window, first use the TAB key to move to that window.

Line status window

Slots 1 and 2 contain the built-in T1 (or E1) lines, with Slot 1 containing the two leftmost lines when you look at the unit's back panel. By default, the top two status windows show the status of the lines in Slot 1:

```
|-----|
|10-100 1234567890 |
| L1/LA nnnnnnnnnn |
| 12345678901234 |
| nnnnn..... |
|-----|
```

Each window displays four lines of information, as follows:

- The first line shows the menu number and column numbers for channels 1–10.
- The second line identifies the line (the Pipeline 220 always indicates L1). It also shows a two-character link status indicator for the line and a one-character status indicator for each channel. For example:
 - LA indicates *Link Active* (the line is physically connected).
 - A lowercase n indicates a channel is nailed.
 - A period (.) indicates a channel is not in use.
- The third line has column headers for channels 11–24.
- The fourth line shows a 1-character channel status indicator for channels 11–24.
 - A lowercase n indicates a channel is nailed.
 - A period (.) indicates a channel is not in use.

System Events

The System Events status window provides a log of up to 32 of the most recent system events the Pipeline 220 has recorded:

```
|-----|
|00-200 11:23:55 |
|>M31 Line Ch |
| LAN session up |
| REMOTEGW |
|-----|
```

The message logs update dynamically. Press the Up Arrow key to display the previous entry. Press the Down Arrow key to display the next entry. The Delete key clears all the messages in the log.

The message log displays the information described in the following paragraphs.

Line 1

The first line of the window shows the status-window number and the time the event occurred.

Line 2

The second line identifies the log entry number (M00–M31) and, if applicable, the line and channel on which the event occurred.

Line 3

The third line includes the text of the message. The message can contain either basic information or a warning.

Line 4

The fourth line includes a message parameter.

Sessions

The Sessions status window indicates the number of active bridging/routing links. An online link, as configured in the Connection Profile, constitutes a single active session. A session can be PPP encapsulated:

```
|-----|
|20-100 Sessions|
|>1 Active      |
| 0 REMOTEGW    |
|               |
|-----|
```

The following paragraphs describe each line of the window.

Line 1

The first line specifies the window number and name of the window.

Line 2

The second line indicates the number of active sessions.

Line 3 and succeeding lines

The third and all remaining lines indicate the state of each active session, and the name, address, or CLID of the remote end. Each line uses the format *y zzzzz*, where *y* is a session status character and *zzzzz* indicates the name, address, or CLID of the remote device.

Table 16-3 lists the session status characters that can appear.

Table 16-3. Session status characters

Character	Description
Blank	No connection exists and no other Pipeline 220 operations are being performed.

Table 16-3. Session status characters (continued)

Character	Description
R	Ringin—an incoming call is ringing on the line, ready to be answered.
A	Answerin—the Pipeline 220 is answering an incoming call.
C	Callin—the Pipeline 220 is dialing an outgoing call.
O	Online—a call is up on the line.
H	Hangin up—the Pipeline 220 is clearing the call.

Dyn Stat

The Dyn Stat window shows the name, quality, bandwidth, and bandwidth utilization of each online connection:

```
|-----|
| 20-500 Dyn Stat |
| Qual Good 00:02:04 |
| OK      0 channels |
| CLU   0%  ALU   0% |
|-----|
```

You can press the Down Arrow key to see other connections. More than one connection can be online at once.

The following paragraphs describe each line of the window.

Line 1

The first line of the Dyn Stat window shows its window number and the name of the current Connection Profile. If no connection is currently active, the window name appears instead.

Line 2

The second line lists the quality of the link and the amount of time the link has been active. When a link is online more than 96 hours, the Pipeline 220 reports the duration in number of days. The link quality can have one of the values listed in Table 16-4.

Table 16-4. Link quality values

Value	Description
Good	The current rate of CRC errors is less than 1%.
Fair	The current rate of CRC errors is between 1% and 5%.

Table 16-4.Link quality values (continued)

Value	Description
Marg	The current rate of CRC errors is between 5% and 10%.
Poor	The current rate of CRC errors is more than 10%.
N/A	The link is not online.

Line 3

The third line of the Dyn Stat window shows the current data rate in Kbps, and how many channels the data rate represents.

Line 4

The fourth line displays the CLU and ALU values.

CLU is the Current Line Utilization: the percentage of bandwidth currently being used by the call, divided by the total amount of bandwidth available.

ALU is the Average Line Utilization: the average amount of available bandwidth used by the call during the current history period as specified by the Sec History and Dyn Alg parameters.

WAN Stat

The WAN Stat window shows the current count of received frames, transmitted frames, and frames with errors for each active WAN link. It also indicates the overall count for all data packets received or transmitted across the WAN:

```
|-----|
| 20-300 WAN Stat |
|>Rx Pkt:  387112 |
| Tx Pkt:   22092 |
|   CRC:    0     |
|-----|
```

The following paragraphs describe each line of the menu.

Line 1

The first line displays the window number and name of the window. You can press the Down Arrow key to get per-link statistics. The first line of a per-link display indicates the name, IP address, or MAC address of the remote device. The per-link count is updated every 30 seconds. The overall count is updated at the end of every active link.

Line 2

The second line specifies the number of received frames.

Line 3

The third line displays the number of transmitted frames.

Line 4

The fourth line indicates the number of errored frames. CRC checking is performed on PPP and MP+ links. An errored CRC frame includes at least one data error.

Ether Stat

The Ether Stat window shows the number of Ethernet frames received and transmitted and the number of collisions at the Ethernet interface:

```
|-----|
|50-400 Ether Stat |
|>Rx Pkt:      106 |
| Tx Pkt:      118 |
|   Col:       0  |
|-----|
```

The window includes the fields described in Table 16-5.

Table 16-5. Ether Stat fields

Field	Contents
Rx Pkt	Number of Ethernet frames received from the Ethernet interface.
Tx Pkt	Number of Ethernet frames transmitted over the Ethernet interface.
Col	Number of collisions detected at the Ethernet interface.

The counts return to zero when the Pipeline 220 is switched off or reset. Otherwise, the counts continuously increase up to the maximum allowed by the display.

Sys Options

The Sys Options window provides a read-only list that identifies your Pipeline 220 and names each of the features with which it has been equipped:

```
|-----|
|00-100 Sys Options |
|>Security Prof:1   |
| Software +5.1Ap4+ |
|S/N:42901          |
|-----|
```

The Sys Options window can contain the information listed in Table 16-6.

Table 16-6.Sys Options information

Option	Description
Security Prof: 1, Security Prof: 2...	Identifies which of the nine Security Profiles as currently in use.
Software	Shows the version and revision of the system ROM code.
S/N	Shows the serial number of the Pipeline 220. The serial number of your Pipeline 220 can also be found on the model-number/serial-number label on the Pipeline 220's bottom panel.
Up: 00:18:02:17	Shows how long since the Pipeline 220 was reset. Time appears in dd:hh:mm:ss format.
Pipeline 220	Shows the Ascend unit is a Pipeline 220. If there are several types of Ascend units at your site, this helps differentiate between units.
Switched Installed or Switched Not Inst	Shows whether the Pipeline 220 can place calls over switched circuits.
Frm Rel Installed or Frm Rel Not Inst	Shows whether or not the Frame Relay option is installed
Sec Acc Installed or Sec Acc Not Inst	Shows whether or not the Secure Access option is installed.
IPsec Installed or IPsec Not Inst	Shows whether or not the IPsec option is installed.
Dyn Bnd Installed or Dyn Bnd Not Inst	Shows whether Dynamic Bandwidth Allocation functionality is available.
ISDN Sig Installed or ISDN Sig Not Inst	Shows whether ISDN signalling is available.

Ether Opt

The Ether Opt window shows the hardware installed on the Pipeline 220:

```
|-----|
|20-700 Ether Opt|
|>Enet I/F: UTP|
|  Adr0: 00c07b6e06f7|
|  Adr1: 00c07b6037b8|
|-----|
```


The window includes the fields described in Table 16-7.

Table 16-7. *Ether Opt* fields

Field	Contents
Enet I/F: UTP	The type of Ethernet connection.
Adr0	MAC Address of the first Ethernet interface of the Pipeline 220.
Addr1	MAC Address of the second Ethernet interface of the Pipeline 220.

Syslog

Syslog is not a Pipeline 220 status display, but an IP protocol that sends system status messages to a host computer, which is known as the Syslog host. This host, specified by the Log Host parameter in the Ethernet Profile, saves the system status messages in a syslog file. These messages are derived from two sources—the Message Log display and the CDR display:



Note: See the UNIX man pages about `logger(1)`, `syslog(3)`, `syslog.conf(5)`, and `syslogd(8)` for detailed information about the syslog daemon. The syslog function requires UDP port 514.

Level 4 (warning) and Level 5 (informational) syslog messages

The Message Log provides the data for level 4 (warning) and level 5 (informational) syslog messages. Level 4 and level 5 messages appear in the following format:

ASCEND: *slot-n port-n / line-n, channel-n, text-1, text-2*

where:

- *slot-n port-n / line-n* is the device address (slot, port or line, and channel). The device address is suppressed when it is not applicable or unknown.
- *text-1* specifies information about the reason for the syslog message. The messages are similar to those shown in the message log window.
- *text-2* specifies the system name, IP address, or MAC address of the remote end of a session for the messages.

Level 5 (notice) syslog messages

The data for level 5 (notice) syslog messages is derived from the CDR display, lines 3 and 4. Level 5 messages appear in the following format:

ASCEND: *call-event-ID event-description slot-n port-n data-svcK phone-n*

where:

- *call-event-ID* specifies the event ID in the CDR display.
- *event-description* is a description of the CDR event.
- *slot-n port-n* is the address of the AIM port, which is not included in the message if the address is not applicable or not known.
- *data-svcK* indicates the data service in use.
- *phone-n* is the phone number.

Examples

Because the Syslog host adds the date, type, and name of all syslog messages from the Pipeline 220, that data is not included in the message format. Some sample syslog entries follow:

```
Oct 21 11:18:07 marcsmax ASCEND: slot 0 port 0, line 1, channel 1,  
No Connection
```

```
Oct 21 11:18:07 marcsmax ASCEND: slot 4 port 1, Call Terminated
```

```
Oct 21 11:19:07 marcsmax ASCEND: slot 4 port 1, Outgoing Call, 123
```

This example shows three messages for the system marcsmax.

Terminal-server command-line interface

The terminal-server command-line interface provides commands for checking routing tables, Frame Relay connections, and other configuration parameters. To access the terminal-server command-line interface, you must have administrative privileges. (See “About Pipeline 220 passwords” on page 16-7).

You can use any of the following methods to open the terminal-server command-line interface:

- From the Main Edit menu, select System > Sys Diag > Term Serv, and press Enter.
- In the Main Edit menu, press Ctrl-D to open the DO menu, and select E=Termsrv.
- Enter the following keystroke sequence (Escape key, left square bracket, Escape key, zero) in rapid succession:

```
Esc [ Esc 0
```

If you have sufficient privileges to access the command line, the Pipeline 220 displays the command-line prompt. For example:

```
** Ascend terminal-server **  
ascend%
```

To display the list of terminal-server commands, either enter a question mark:

```
ascend% ?
```

or the Help command:

```
ascend% help
```

The system responds by listing the terminal-server commands, with brief explanations:

?	Display help information
help	" " "
quit	Closes terminal-server session
hangup	" " " "
test	test <phone-number> [<frame-count>]
local	Go to local mode
remote	remote <station>
set	Set various items. Type 'set ?' for help
show	Show various tables. Type 'show ?' for help
iproute	Manage IP routes. Type 'iproute ?' for help
dnstab	Manage local DNS table. Type 'dnstab ?' for help
slip	SLIP command
cslip	Compressed SLIP command
ppp	PPP command
menu	Host menu interface
telnet	telnet [-a -b -t] <host-name> [<port-number>]
tcp	tcp <host-name> <port-number>
ping	ping <host-name>
ipxping	ipxping <server-name>
traceroute	Trace route to host. Type 'traceroute -?' for help
kill	kill <session ID>

Exiting the terminal server interface

The following commands close the terminal-server command-line interface and return the cursor to the VT100 menus.

- Quit
- Hangup
- Local

Commands not supported on the Pipeline 220

The Pipeline 220 does not support the following terminal-server commands:

- Test
- Remote
- Slip
- CSlip
- PPP
- Show ISDN
- Show Pools
- Set Password

Commands for use by terminal-server users

The commands described in this section initiate a session with a host or toggle to a different interface that displays a menu selection of Telnet hosts.

Set command

The Set command takes several arguments. To display them, enter the Set command with a question mark:

```
ascend% set ?  
  
set ?           Display help information  
set all         Display current settings  
set term        Sets the telnet/rlogin terminal type  
set password    Enable dynamic password serving  
set circuit     Frame Relay Circuit control
```

Set All and Set Term

The Set All command displays current settings. For example:

```
ascend% set all  
  
term = VT100  
dynamic password serving = disabled
```

To specify a terminal type other than the default VT100, use the Set Term command.

Set Circuit

The Set Circuit command enables you to turn off traffic going through a Frame Relay circuit without disabling the circuit endpoints. This command prevents traffic from going between endpoints without disrupting the state of the DLCI. To display the support options, enter the Set Circuit command with a question mark:

```
ascend% set circuit ?  
  
set circuit ?      Display help information  
set circuit active [name] Set the CIRCUIT to active  
set circuit inactive [name] Set the CIRCUIT to inactive
```

To allow data to flow through a circuit, use the active parameter. For example:

```
ascend% set circuit active circuit-1
```

To turn off data flow without disrupting the state of the DLCIs, use the inactive parameter. For example:

```
ascend% set circuit inactive circuit-2
```

Show command

The Show command takes several arguments. To display them, enter the Show command with a question mark:

```
ascend% show ?
```



```

show ?          Display help information
show arp        Display the Arp Cache
show icmp       Display ICMP information
show if         Display Interface info. Type 'show if ?' for help.
show ip         Display IP information. Type 'show ip ?' for help.
show udp        Display UDP information. Type 'show udp ?' for help.
show igmp       Display IGMP information. Type 'show igmp ?' for help.
show mROUTING   Display MROUTING information.Type 'show mROUTING ?'
show ospf       Display OSPF information. Type 'show ospf ?' for help.
show tcp        Display TCP information. Type 'show tcp ?' for help.
show dnstab     Display local DNS table. Type 'show dnstab ?' for help.
show netware    Display IPX information. Type 'show netware ? ' for
show isdn       Display ISDN events.  Type 'show isdn <line number>'
show fr         Display Frame relay info. Type 'show fr ?' for help.
show pools      Display the assign address pools.
show uptime     Display system uptime.
show revision   Display system revision.
show users      Display concise list of active users
show sessid     Display current and base session id

```

Note: Not all displayed Show commands apply to the Pipeline 220. See this section for specific information.

Displaying the ARP cache

To display the ARP cache, enter the Show ARP command. For example:

```
ascend% show arp
```

entry	typ	ip address	ether addr	if	rtr	pkt	insert
0	DYN	10.65.212.199	00C07B605C07	0	0	0	857783
1	DYN	10.65.212.91	0080C7C4CB80	0	0	0	857866
2	DYN	10.65.212.22	080020792B4C	0	0	0	857937
3	DYN	10.65.212.3	0000813DF048	0	0	0	857566
4	DYN	10.65.212.250	0020AFF80F1D	0	0	0	857883
5	DYN	10.65.212.16	0020AFEC0AFB	0	0	0	857861
6	DYN	10.65.212.227	00C07B5F14B6	0	0	0	857479
7	DYN	10.65.212.36	00C07B5E9AA5	0	0	0	857602
8	DYN	10.65.212.71	0080C730041F	0	0	0	857721
9	DYN	10.65.212.5	0003C6010512	0	0	0	857602
10	DYN	10.65.212.241	0080C72ED212	0	0	0	857781
11	DYN	10.65.212.120	0080C7152582	0	0	0	857604
12	DYN	10.65.212.156	0080A30ECE6D	0	0	0	857901
13	DYN	10.65.212.100	00C07B60E28D	0	0	0	857934
14	DYN	10.65.212.1	00000C065D27	0	0	0	857854
15	DYN	10.65.212.102	08000716C449	0	0	0	857724
16	DYN	10.65.212.33	00A024AA0283	0	0	0	857699
17	DYN	10.65.212.96	0080C7301792	0	0	0	857757
18	DYN	10.65.212.121	0080C79BF681	0	0	0	857848
19	DYN	10.65.212.89	00A024A9FB99	0	0	0	857790
20	DYN	10.65.212.26	00A024A8122C	0	0	0	857861
21	DYN	10.65.212.6	0800207956A2	0	0	0	857918
22	DYN	10.65.212.191	0080C75BE778	0	0	0	857918
23	DYN	10.65.212.116	0080C72F66CC	0	0	0	857416

```
24 DYN 10.65.212.87 0000813606A0 0 0 0 857666
25 DYN 10.65.212.235 00C07B76D119 0 0 0 857708
26 DYN 10.65.212.19 08002075806B 0 0 0 857929
```

In the output:

- `entry` is a unique identifier for each ARP table entry.
- `typ` specifies how the address was learned, dynamically (DYN) or statically (STAT).
- `ip Address` indicates the address contained in ARP requests.
- `ether Addr` indicates the MAC address of the host with that IP address.
- `if` specifies the interface on which the Pipeline 220 received the ARP request.
- `rtr` is the next-hop router on the specified interface.

Displaying ICMP packet statistics

To view the number of ICMP packets received intact, received with errors, and transmitted, enter the Show ICMP command. For example:

```
ascend% show icmp
3857661 packet received.
20 packets received with errors.
    Input histogram: 15070
2758129 packets transmitted.
0 packets transmitted due to lack of resources.
    Output histogram: 15218
```

The Input and Output histograms show the number of ICMP packets received and transmitted in each category.

Displaying interface statistics

To display the supported interface commands, enter the Show IF command with a question mark:

```
ascend% show if ?
show if ?          Display help information
show if stats      Display Interface Statistics
show if totals     Display Interface Total counts
```

To display the status and packet count of each active WAN link and of the local and loopback interfaces, enter the Show IF Stats command. For example:

```
ascend% show if stats
```

Interface	Name	Status	Type	Speed	MTU	InPackets	Oupackets
ie0	ethernet	Up	6	10000000	1500	107385	85384
wan0		Down	1	0	1500	0	0
wan1		Down	1	0	1500	0	0
wan2		Down	1	0	1500	0	0
wanidle0		Up	6	10000000	1500	0	0
lo0	loopback	Up	24	1000000	1500	0	0

In the output:

- Interface specifies the interface name (see the *Pipeline 220 Interface Configuration Guide*)
- Name is the name of the profile or a text name for the interface
- Status indicates either Up (the interface is functional), or Down.
- Type specifies the type of application being used on the interface, as specified in RFC 1213 (MIB-2). For example, 23 indicates PPP and 28 indicates SLIP.
- Speed is the data rate in bits per second.
- MTU is the maximum packet size allowed on the interface. MTU stands for Maximum Transmission Unit.
- InPacket is the number of packets the interface has received.
- OutPackets is the number of packets the interface has transmitted.

To display the packet count at each interface, broken down by type of packet, enter the Show IF Totals command. For example:

```
ascend% show if totals
```

Name	--Octets----	Ucast--	-NonUcast-	Discard	-Error-	Unknown	-Same	IF-
ie0 i:	7813606	85121	22383	0	0	0		0
o:	101529978	85306	149	0	0	0		0
wan0 i:	0	0	0	0	0	0		0
o:	0	0	0	0	0	0		0
wan1 i:	0	0	0	0	0	0		0
o:	0	0	0	0	0	0		0
wan2 i:	0	0	0	0	0	0		0
o:	0	0	0	0	0	0		0
wanidle0 i:	0	0	0	0	0	0		0
o:	0	0	0	0	0	0		0
lo0 i:	0	0	0	0	0	0		0
o:	0	0	0	0	0	0		0

In the output:

- Name is the interface name (see the *Pipeline 220 Interface Configuration Guide*)
- Octets is the total number of bytes processed by the interface.
- Ucast is the number of packets with a unicast destination address.
- NonUcast is the number of packets with a multicast address or a broadcast address.
- Discard is the number of packets that the interface could not process.
- Error is the number of packets with CRC errors, header errors, or collisions.
- Unknown is the number of packets the Pipeline 220 forwarded across all bridged interfaces because of unknown or unlearned destinations.
- Same IF is the number of bridged packets whose destination is the same as the source.

Displaying IP statistics and addresses

To display the supported IP commands, enter the Show IP command with a question mark:

```
ascend% show ip ?
```

```
show ip ?          Display help information
show ip stats      Display IP Statistics
```

```
show ip address      Display IP Address Assignments
show ip routes       Display IP Routes
```

To display statistics about IP activity, including the number of IP packets the Pipeline 220 has received and transmitted, enter the Show IP Stats command. For example:

```
ascend% show ip stats
107408 packets received.
    0 packets received with header errors.
    0 packets received with address errors.
    0 packets forwarded.
    0 packets received with unknown protocols.
    0 inbound packets discarded.
107408 packets delivered to upper layers.
85421 transmit requests.
    0 discarded transmit packets.
    1 outbound packets with no route.
    0 reassembly timeouts.
    0 reassemblies required.
    0 reassemblies that went OK.
    0 reassemblies that Failed.
    0 packets fragmented OK.
    0 fragmentations that failed.
    0 fragment packets created.
    0 route discards due to lack of memory.
64 default ttl.
```

To view IP interface address information, enter the Show IP Address command. For example:

```
ascend% show ip address
```

Interface	IP Address	Dest Address	Netmask	MTU	Status
ie0	10.2.3.4	N/A	255.255.255.224	1500	Up
wan0	0.0.0.0	N/A	0.0.0.0	1500	Down
wan1	13.1.2.0	13.1.2.128	255.255.255.248	1500	Down
wan2	0.0.0.0	N/A	0.0.0.0	1500	Down
wan3	0.0.0.0	N/A	0.0.0.0	1500	Down
lo0	127.0.0.1	N/A	255.255.255.255	1500	Up
rj0	127.0.0.2	N/A	255.255.255.255	1500	Up
bh0	127.0.0.3	N/A	255.255.255.255	1500	Up

Displaying UDP statistics and listen table

To display the supported UDP commands, enter the Show UDP command with a question mark:

```
ascend% show udp ?
show udp ?          Display help information
show udp stats      Display UDP Statistics
show udp listen     Display UDP Listen Table
```

To display the number of UDP packets received and transmitted, enter the Show UDP Stats command. For example:

```
ascend% show udp stats
```

```
22386 packets received.
    0 packets received with no ports.
    0 packets received with errors.
    0 packets dropped
    9 packets transmitted.
```

To view information about the socket number, UDP port number, and the number of packets queued for each UDP port on which the Pipeline 220 is currently listening, enter the Show UDP Listen command. For example:

```
ascend% show udp listen
```

Socket	Local Port	InQLen
0	520	0
1	7	0
2	123	0
3	514	0
4	161	0
5	162	0

Viewing multicast interfaces

For viewing multicast interfaces, the Pipeline 220 supports Internet Group Management Protocol (IGMP) commands and multicast routing (mrouting) commands.

To display the supported IGMP commands, enter the Show IGMP command with a question mark:

```
ascend% show igmp ?
```

```
show igmp ?          Display help information
show igmp stats      Display IGMP Statistics
show igmp groups     Display IGMP groups Table
show igmp clients    Display IGMP clients
```

To display the mrouting commands, enter the Show mrouting command with a question mark:

```
ascend% show mrouting ?
```

```
show mrouting ?      Display help information
show mrouting stats  Display MROUTING Statistics
```

Displaying the multicast forwarding table

To display active multicast group addresses and clients (interfaces) registered for each group, enter the Show IGMP Groups command. For example:

```
ascend% show igmp groups
```

```
IGMP Group address Routing Table Up Time: 0:0:22:17
Hash      Group Address  Members  Expire time  Counts
  10      224.0.2.250
                2          0:3:24      3211 :: 0 S5
                1          0:3:21      145 :: 0 S5
                0 (Mbone)  .....    31901 :: 0 S5
```

In the output:

- Hash is an index to a hash table (displayed for debugging purposes only).

- `Group` address indicates the IP multicast address used in this packet.

Note: The IP multicast address being monitored is marked with an asterisk, meaning that this address is joined by local application.

- `Members` is the interface ID on which the membership resides. 0 represents the Ethernet interface. Other numbers represent WAN interfaces, numbered according to when they became active. The interface labeled `Mbone` is the interface on which the multicast router resides.
- `Expire time` indicates when this membership expires. The Pipeline 220 sends out IGMP queries every 60 seconds, so the expiration time is usually renewed. If the expiration time is reached, the entry is removed from the table. Periods in this field indicate that this membership never expires.
- `Counts` is the number of packets forwarded to the client, the number of packets dropped due to lack of resources, and the state of the membership (the state appears for debugging purposes).

Listing multicast clients

To display a list of multicast clients, enter the `Show IGMP Clients` command. For example:

```
ascend% show igmp clients
```

IGMP Clients

Client	Version	RecvCount	CLU	ALU
0 (Mbone)	1	0	0	0
2	1	39	68	67
1	1	33310	65	65

In the output:

- `Client` indicates the interface ID on which the client resides. 0 represents the Ethernet. Other numbers are WAN interfaces, numbered according to when they became active. The interface labeled `Mbone` is the interface on which the multicast router resides.
- `Version` is the version of IGMP being used.
- `RecvCount` is the number of IGMP messages received on that interface.
- `CLU` (Current Line Utilization) indicates percentage of bandwidth currently utilized across the interface. If bandwidth utilization is high, some IGMP packet types will not be forwarded.
- `ALU` (Average Line Utilization) indicates percentage of bandwidth utilized across the interface. If bandwidth utilization is high, some IGMP packet types will not be forwarded.

Displaying multicast activity

To display the number of IGMP packet types sent and received, enter the `Show IGMP Stats` command. For example:

```
ascend% show igmp stats
```

```
46 packets received.
0 bad checksum packets received.
0 bad version packets received.
0 query packets received.
```

```
46 response packets received.  
0 leave packets received.  
51 packets transmitted.  
47 query packets sent.  
4 response packets sent.  
0 leave packets sent.
```

Viewing Multicast Routing statistics

To display the number of multicast packets received and forwarded, enter the Show MRouting Stats command. For example:

```
ascend% show mrouting stats  
34988 packets received.  
57040 packets forwarded.  
0 packets in error.  
91 packets dropped.  
0 packets transmitted.
```

In many cases, the number of packets forwarded will be greater than the number of packets received, because packets can be duplicated and forwarded across multiple links.

Monitoring OSPF

To display the supported commands, enter the Show OSPF command with a question mark:

```
ascend% show ospf ?  
  
show ospf ?                Display help information  
show ospf errors            Display OSPF errors  
show ospf areas             Display OSPF areas  
show ospf general           Display OSPF general info  
show ospf interfaces        Display OSPF interfaces  
show ospf lsdb              Display OSPF link-state DB  
show ospf lsa               Display OSPF link-state advertisements  
show ospf nbrs              Display OSPF neighbors  
show ospf rtab              Display OSPF routing tab  
show ospf io                Display OSPF io
```

Viewing OSPF errors

To view OSPF errors, enter the Show OSPF Errors command. For example:

```
ascend% show ospf errors  
  
ERRORS from:                boot  
0: IP: Bad OSPF pkt type      0: IP: Bad IP Dest  
0: IP: Bad IP proto id       1: IP: Pkt src = my IP addr  
0: OSPF: Bad OSPF version     0: OSPF: Bad OSPF checksum  
0: OSPF: Bad intf area id     0: OSPF: Area mismatch  
0: OSPF: Bad virt link info   0: OSPF: Auth type != area type  
0: OSPF: Auth key != area key 0: OSPF: Packet is too small  
0: OSPF: Packet size > IP length 0: OSPF: Transmit bad  
0: OSPF: Received on down IF  0: Hello: IF mask mismatch  
0: Hello: Unknown Virt nbr    0: Hello: Unknown NBMA nbr
```

```

0: DD: Unknown nbr                0: DD: Nbr state low
0: DD: Nbr's rtr = my rtrid       0: DD: Extern option mismatch
0: Ack: Unknown nbr               0: Ack: Nbr state low
0: Ls Req: Nbr state low           0: Ls Req: Unknown nbr
0: Ls Req: Empty request           0: Ls Req: Bad pkt
0: LS Update: Nbr state low        0: Ls Update: Unknown nbr
0: Ls Update: Newer self-gen LSA   0: Ls Update: Bad LS chksum

```

The output lists all error messages related to OSPF, with each message preceded by the number of times it has been generated since the Pipeline 220 powered up. Immediately following the number is a field indicating the packet type:

- IP (IP packets)
- OSPF (OSPF packets)
- Hello (Hello packets)
- DD (Database Description packets, which are exchanged periodically between neighbors)
- Ack (every DD packet must be acknowledged)
- LS Req (Link-state request— a request for an updated database)
- LS Update (An exchange to update databases)

Viewing OSPF areas

To view information about OSPF areas, enter the Show OSPF Areas command. For example:

```

ascend% show ospf areas
Area ID: 0.0.0.0
Auth Type: Simple Passwd  Import ASE: On  Spf Runs: 23
Local ABRs: 0  Local ASBRs: 5  Inter LSAs: 7  Inter Cksum sum: 0x2ee0e

```

In the output:

- `Area ID` specifies the area number in dotted-decimal format.
- `Auth Type` indicates the type of authentication, Simple or Null.
- `Import ASE` specifies route calculation method. In effect, it specifies whether the router is an ABR or not. This functionality is always ON in the Pipeline 220.
- `Spf Runs` is the number of times the SPF calculation was run. The calculation is performed every time the router notes a topology change or receives an update from another router.
- `Local ABRs` is the number of ABRs the router knows about and the number of areas. Zero (0) indicates the router knows about the backbone area only.
- `Local ASBRs` is the number of ASBRs the router knows about.
- `Inter LSAs` shows the number of entries in the link-state database.
- `Inter Cksum sum` shows the checksum to indicate if a database has changed.

Viewing OSPF general info

To display general information about OSPF, enter the Show OSPF General command. For example:

```

ascend% show ospf general

```



```
Rtr ID: 10.5.2.154
Status: Enabled Version: 2 ABR: Off ASBR: On
LS ASE Count: 8 ASE Cksum sum: 0x4c303 Tos Support: TOS 0 Only
New LSA Originate Count: 13 Rx New LSA Count: 498
```

In the output:

- `Rtr ID` is the IP address assigned to the Pipeline 220 Ethernet interface.
- `Status` shows whether OSPF is enabled or disabled.
- `Version` is the version of the OSPF protocols running.
- `ABR` can be On or Off, depending on where the Pipeline 220 is situated on the network. If ABR is on, the Pipeline 220 performs additional calculations related to external routes.
- `ASBR` is always displayed as On in the Pipeline 220. Although the Pipeline 220 cannot function as an IGP gateway, it does import external routes (for example, when it establishes a WAN link with a caller that does not support OSPF) and the ASBR calculations are always performed.
- `LS ASE count` is the number of link-state database entries that are external.
- `ASE Cksum sum` specifies a checksum used to note that ASE routes in the database have changed.
- `TOS Support` indicates the level of TOS support in the router.
- `New LSA Originate Count` is the number of LSAs this router created.
- `Rx New LSA Count` is the number of LSAs this router received from other OSPF routers.

To display the OSPF interfaces, enter the Show OSPF Interfaces command. For example:

```
ascend% show ospf interfaces
```

Area	IP Address	Type	State	Cost	Pri	DR	BDR
0.0.0.0	10.5.2.154	Bcast	BackupDR	1	5	10.5.2.155	10.5.2.154
0.0.0.0	10.5.2.154	PtoP	P To P	10	5	None	None
0.0.0.0	10.5.2.154	PtoP	P To P	10	5	None	None

In the output:

- `Area` shows the area ID (0.0.0.0 is the backbone).
- `IP Address` is the address assigned to the Pipeline 220's Ethernet interface. To identify WAN links, use the Type and Cost fields.
- `Type` can be broadcast or point-to-point. WAN links are point-to-point.
- `State` shows how far along the router is in the election process of a DR or BDR. The state can be 1-way (indicating that the election process has begun), 2-way (indicating that the router has received notification), BackupDR, or DR.
- `Cost` is the metric assigned to the link. The default cost for Ethernet is 1.
- `Pri` shows the designated router election priority assigned to the Pipeline 220.
- `DR` identifies the designated router.
- `BDR` identifies the backup designated router.

Viewing the OSPF link-state database

To view the router's link-state database, enter the Show OSPF LSDB command. For example:

```
ascend% show ospf lsdb
```

LS Data Base:

Area	LS Type	Link ID	Adv Rtr	Age	Len	Seq #	Metric
0.0.0.0	STUB	10.5.2.146	10.5.2.146	3600	24	0	0
0.0.0.0	STUB	10.5.2.154	10.5.2.154	3600	24	0	0
0.0.0.0	STUB	10.5.2.155	10.5.2.155	3600	24	0	0
0.0.0.0	STUB	10.5.2.162	10.5.2.162	3600	24	0	0
0.0.0.0	STUB	10.5.2.163	10.5.2.163	3600	24	0	0
0.0.0.0	RTR	10.5.2.146	10.5.2.146	659	72	8000003e	0
0.0.0.0	RTR	10.5.2.154	10.5.2.154	950	84	8000000a	0
0.0.0.0	RTR	10.5.2.155	10.5.2.155	940	60	80000005	0
0.0.0.0	RTR	10.5.2.162	10.5.2.162	980	84	8000003b	0
0.0.0.0	RTR	10.5.2.163	10.5.2.163	961	60	80000005	0
0.0.0.0	NET	10.5.2.155	10.5.2.155	940	32	80000003	0
0.0.0.0	NET	10.5.2.163	10.5.2.163	961	32	80000003	0
0.0.0.0	ASE	10.5.2.16	10.5.2.163	18	36	80000098	3
0.0.0.0	ASE	10.5.2.18	10.5.2.163	546	36	80000004	10
0.0.0.0	ASE	10.5.2.144	10.5.2.146	245	36	80000037	1
0.0.0.0	ASE	10.5.2.152	10.5.2.154	536	36	80000006	1
0.0.0.0	ASE	10.5.2.152	10.5.2.155	526	36	80000004	1
0.0.0.0	ASE	10.5.2.152	10.5.2.163	18	36	80000097	9
0.0.0.0	ASE	10.5.2.155	10.5.2.163	17	36	80000097	9
0.0.0.0	ASE	10.5.2.160	10.5.2.162	568	36	80000037	1

In the output:

- **Area** specifies the area ID.
- **LS Type** indicates the type of link as defined in RFC 1583:
 - Type 1 (RTR) are router-LSAs that describe the collected states of the router's interfaces.
 - Type 2 (NET) are network-LSAs that describe the set of routers attached to the network.
 - Types 3 and 4 (STUB) are summary-LSAs that describe point-to-point routes to networks or AS boundary routers.
 - Type 5 (ASE) are AS-external-LSAs that describe routes to destinations external to the Autonomous System. A default route for the Autonomous System can also be described by an AS-external-LSA.
- **Link ID** is the target address of the route.
- **Adv Rtr** is the address of the advertising router.
- **Age** is the age of the route in seconds.
- **Len** is the length of the LSA.
- **Seq #** is a number that begins with 80000000 and increments by one for each LSA received.
- **Metric** is the cost of the link, not of a route. The cost of a route is the sum of all intervening links, including the cost of the connected route.

You can expand each entry in the link-state database to view additional information about a particular LSA.

Viewing OSPF link-state advertisements

To specify a link-state advertisement to be expanded, use the following format for the Show OSPF LSA command:

```
show ospf lsa area ls-type ls-id adv-rtr
```

The command requires that you include the first four fields of the LSA as listed in the database. You can select the first four fields and paste them in after typing the command. For example, to show an expanded view of the last entry in the link-state database shown in the previous section:

```
ascend% show ospf lsa 0.0.0.0 ase 10.5.2.160 10.5.2.162
LSA  type: ASE ls id: 10.5.2.160 adv rtr: 110.5.2.162 age: 568
      len: 36 seq #: 80000037 cksum: 0xffff
      Net mask: 255.255.255.255 Tos 0 metric: 10 E type: 1
      Forwarding Address: 0.0.0.0 Tag: c0000000
```

Viewing OSPF neighbors

To view adjacencies, enter the Show OSPF NBRs command. For example:

```
ascend% show ospf nbrs
```

Area	Interface	Router Id	Nbr IP Addr	State	Mode	Pri
0.0.0.0	10.5.2.154	10.5.2.155	10.5.2.155	Full	Slave	5
0.0.0.0	10.5.2.154	10.5.2.146	10.5.2.146	Full	Master	5
0.0.0.0	10.5.2.154	10.5.2.162	10.5.2.162	Full	Slave	5

In the output:

- **Area** is the area ID.
- **Interface** shows the address assigned to the interface. In the Pipeline 220, the IP address is always the address assigned to the Ethernet interface.
- **Router Id** is the IP address of the router used to reach a neighbor. This is often the same address as the neighbor itself.
- **Nbr IP Addr** is the IP address of the neighbor.
- **State** indicates the state of the link-state database exchange. Full means that the databases are fully aligned between the Pipeline 220 and its neighbor.
- **Mode** indicates whether the neighbor is functioning in master or slave mode. The master sends Database Description packets (polls), which are acknowledged by Database Description packets sent by the slave (responses).
- **Pri** indicates the designated router election priority assigned to the Pipeline 220.

Viewing the OSPF routing table

To view the OSPF routing table, enter the Show OSPF Rtab command. For example:

```
ascend% show ospf rtab
```

SPF algorithm	run	24 times	since	boot	
Dest	D_mask	Area	Cost E Path	Nexthop	AdvRtr

Nets:					

```

10.5.2.163    255.255.255.248    0.0.0.0    10    3    EXT    10.5.2.163    10.5.2.16
10.5.2.163    255.255.255.255    0.0.0.0    20    0    EXT    10.5.2.163    10.5.2.16
10.5.2.146    255.255.255.248    0.0.0.0    20    1    EXT    10.5.2.154    10.5.2.14
10.5.2.146    255.255.255.255    0.0.0.0    20    0    STUB    10.5.2.154    10.5.2.14
10.5.2.155    255.255.255.248    0.0.0.0    10    0    INT    10.5.2.154    10.5.2.15
10.5.2.154    255.255.255.255    0.0.0.0    21    0    STUB    10.5.2.163    10.5.2.15
10.5.2.155    255.255.255.255    0.0.0.0    20    9    STUB    10.5.2.155    10.5.2.15
10.5.2.163    255.255.255.248    0.0.0.0    11    1    INT    10.5.2.163    10.5.2.16
10.5.2.162    255.255.255.255    0.0.0.0    20    0    STUB    10.5.2.163    10.5.2.16
10.5.2.163    255.255.255.255    0.0.0.0    10    0    STUB    10.5.2.163    10.5.2.16

```

In the output:

- `Dest` shows the destination address.
- `D_mask` is the destination netmask.
- `Area` is the area ID.
- `Cost` is the cost of the route.
- `E` is the cost of the link. (The cost of a route is the sum of the cost of each intervening link, including the cost to the connected route.)
- `Path` specifies the type of link: EXT (exterior), INT (interior), or STUB (a default).
- `Next hop` specifies the target address from this router.
- `Adv Rtr` is the advertising router. Sometimes a router will advertise routes for which it is not the gateway.

Viewing OSPF protocol I/O

To display information about packets sent and received by the OSPF protocol, enter the `Show OSPF IO` command. For example:

```

ascend% show ospf io

IO stats from:                                boot
>> RECEIVED:
    0: Monitor request
   785: Hello
    13: DB Description
     6: Link-State Req
  1387: Link-State Update
    64: Link-State Ack
>> SENT:
   794: Hello
    15: DB Description
     6: Link-State Req
  1017: Link-State Update
   212: Link-State Ack

```

Displaying TCP statistics and connections

To display the commands available for showing TCP statistics and connections, enter the `Show TCP` command with a question mark:

```

ascend% show tcp ?

```

```
show tcp ?           Display help information
show tcp stats       Display TCP Statistics
show tcp connection Display TCP Connection Table
```

To display the number of TCP packets received and transmitted, enter the Show TCP Stats command. For example:

```
ascend% show tcp stats
      0 active opens.
     11 passive opens.
      1 connect attempts failed.
      1 connections were reset.
      3 connections currently established.
  85262 segments received.
  85598 segments transmitted.
    559 segments re-transmitted.
```

An active open is a TCP session that the Pipeline 220 initiated. A passive open is a TCP session that the Pipeline 220 did not initiate.

To display current TCP sessions, enter the Show TCP Connections command. For example:

```
ascend% show tcp connection
```

Socket	Local	Remote	State
0	*.23	*,*	LISTEN
1	10.2.3.23	15.5.248.121.15003	ESTABLISHED

Displaying IPX packet statistics

To display IPX packet statistics, enter the Show NetWare Stats command. For example:

```
ascend% show netware stats
27162 packets received.
25392 packets forwarded.
0 packets dropped exceeding maximum hop count.
0 outbound packets with no route.
```

The Pipeline 220 drops packets that exceed the maximum hop count (that have already passed through too many routers).

Displaying the IPX service table

To display the IPX service table, enter the Show NetWare Servers command. For example:

```
ascend% show netware servers
```

IPX address	type	server name
ee000001:000000000001:0040	0451	server-1

In the output:

- `IPX address` is the address of the server. The address uses the following format:
network number:node number:socket number
- `type` indicates the type of service available (in hexadecimal format). For example, 0451 designates a file server.
- `server name` is the first 35 characters of the server name.

Displaying the IPX routing table

To display the IPX routing table, enter the Show NetWare Networks command. For example:

```
ascend% show netware networks

network      next router      hops      ticks      origin
CFFF0001     000000000000         0          1      Ethernet      S
```

The output includes the following fields:

- `network`: The IPX network number.
- `next router`: The address of the next router, or 0 (zero) for a direct or WAN connection.
- `hops`: The hop count to the network.
- `ticks`: The tick count to the network.
- `origin`: The name of the profile used to reach the network.

Note: An S or an H flag can appear next to the origin. S indicates a static route. H indicates a hidden static route. Hidden static routes occur when the router learns of a better route.

Monitoring Frame Relay connections

To display the commands available for showing Frame Relay statistics and connections, enter the Show FR command with a question mark:

```
ascend% show fr ?

show fr ?          Display help information
show fr stats      Display Frame relay information
show fr lmi        Display Frame relay LMI information
show fr dlci [name] Display all DLCI information or just for [name]
show fr circuits   Display the FR Circuit table
```

Displaying Frame Relay statistics

To display Frame Relay statistics, enter the Show FR Stats command. For example:

```
ascend% show fr stats

Name      Type   Status   Speed   MTU   InFrame   OutFrame
fr1       DCE    Down     64000   1532   0         1
fr1-temp  DCE    Up       64000   1532   0         1
fr1-temp-9 DCE    Up       64000   1532   0         0
```

In the output:

- `Name` is the name of the Frame Relay profile associated with the interface.
- `Type` indicates the type of interface.
- `Status` indicates the status of the interface. Up shows the interface is functional, but is not necessarily handling an active call. Down shows the interface is not functional.
- `Speed` is the data rate in bits per second.
- `MTU` is the maximum packet size allowed on the interface.
- `InFrame` is the number of frames the interface has received.
- `OutFrame` is the number of frames transmitted.

Displaying link management information

To display Link Management Information (LMI) for each link activated by a Frame Relay profile, enter the Show FR LMI command. For example:

```
ascend% show fr lmi
T1_617D LMI for fr1
  Invalid Unnumbered info      0  Invalid Prot Disc      0
  Invalid Dummy Call Ref      0  Invalid Msg Type      0
  Invalid Status Message      0  Invalid Lock Shift    0
  Invalid Information ID       0  Invalid Report Type   0
  Num Status Enqs Sent        0  Num Status Msgs Rcvd  0
  Num Update Status Rcvd      0  Num Status Timeouts   2779
LMI is not on for fr1-temp
LMI is not on for fr1-temp-9
```

This information is based on the ANSI T1.617 Annex D local in-channel signaling protocol. (See Annex D for a full definition of each of the fields reported.)

Displaying DLCI status

To display the status of each DLCI, enter the Show FR DLCI command. For example:

```
ascend% show fr dlci
DLCIs for fr1
DLCIs for fr1-temp
eng-lab-236-Cir   DLCI =   17   Status = ACTIVE
  input pkts      0          output pkts      0
  input octets    0          output octets    0
  input FECN      0          input DE          0
  input BECN      0
last time status changed: 03/05/1997 14:44:17
DLCIs for fr1-temp-9
eng-lab-236-Cir-9 DLCI =   16   Status = ACTIVE
  input pkts      0          output pkts      0
  input octets    0          output octets    0
  input FECN      0          input DE          0
  input BECN      0
last time status changed: 03/05/1997 14:45:07
DLCIs not assigned
```

In the output:

- DLCI is the DLCI number.
- Status indicates ACTIVE if the connection is up or INACTIVE if not.
- input pkts is the number of frames the interface has received.
- output pkts is the number of frames the interface has transmitted.
- input octets is the number of bytes the interface has received.
- output octets is the number of bytes the interface has transmitted.
- input FECN is the number of packets received with the FECN (Forward Explicit Congestion Notification) bit set. This field always includes a 0 (zero) because congestion management is not currently supported.

- `input BECN` is the number of packets received with the BECN (Backward Explicit Congestion Notification) bit set. This field always includes a 0 (zero) because congestion management is not currently supported.
- `DE Pkts` is the number of packets received with the DE (Discard Eligibility) indicator bit set.
- `Last Time Status Changed` indicates the last time the DLCI state changed.

Displaying circuit information

To display the Frame Relay profile name, DLCI, and status of configured circuits, enter the `Show FR Circuits` command. For example:

```
ascend% show fr circuits
cir-9 User Setting Up
fr1-temp-9          16          Up
fr1-temp            17          Up
```

Show Uptime

To view how long the Pipeline 220 has been running, enter the `Show Uptime` command. For example:

```
ascend% show uptime
system uptime: up 2 days, 4 hours, 38 minutes, 43 seconds
```

If the Pipeline 220 stays up for 1000 consecutive days with no power cycles, the number of days displayed *turns over* to 0 and begins to increment again.

Show Revision

To display load and version number of the software currently running in the Pipeline 220, enter the `Show Revision` command. For example:

```
ascend% show revision
techpubs-lab-17 system revision: ebiom.m40 5.0A
```

Show Users

To display the number of active sessions:

```
ascend% show users
```

I	Session	Line:	Slot:	Tx	Rx	Service	Host	User
O	ID	Chan	Port	Data	Rate	Type[mpID]	Address	Name
O	245761821	32459	2:1	n/a	n/a	Frame relay	10.10.212.10	corp

In the output:

- `IO` indicates either I (incoming call) or O (outgoing call).
- `Session ID` is the unique session-ID.
- `Line:Chan` shows the line and channel of the established session.
- `Slot:Port` shows the slot and port of the service being used by the session. The slot data can indicate the number of a slot containing a modem card and the modem on that

card, or the virtual slot of the Pipeline 220 unit's bridge/router. The port data shows the virtual interfaces to the bridge/router, starting with 1 for the first session of a multichannel session.

- **Data Rate** indicates the bearer capacity or modem speed, as appropriate to the session type.
- **Service Type** specifies the type of session, either **Termsrv** or a protocol name. The special values **Initial** and **Login** document the progress of a session. **Initial** identifies sessions that do not yet have a protocol assigned.
- **Host Address** shows the network address of the host originating the session. In some cases, this field might be **N/A**.
- **User Name** specifies the station name associated with the session. Initially, this value is **Answer**. This is usually replaced with the name of the remote host.

Show SessID

The **Show SessID** command displays the current internal session identification number available for the Pipeline 220 to assign to the next connection. At a given time, the Pipeline 220 has assigned the saved base value to the first connection after its last reboot. Subsequent connections are assigned new session numbers, incremented one from the previous session number. For example:

```
ascend% show sessid
Session ID current 243975689, saved base 243975685
```

IProute command

The terminal-server **IProute** commands display the routing table and enable you to add or delete routes. The changes you make to the routing table by using the **IProute** command last only until the Pipeline 220 unit resets. To view the supported commands, enter the **IProute** command with a question mark:

```
ascend% iproute ?

iproute ?      Display help information
iproute add    iproute add <destination/size> <gateway> [ pref ] [ m ]
iproute delete iproute delete <destination/size> <gateway> [ proto ]
iproute show   displays IP routes (same as "show ip routes" command)
```

Displaying the routing table

Note that the **IProute Show** command and the **Show IP Routes** command have identical output. To view the IP routing table, enter the **IProute Show** command. For example:

```
ascend% iproute show
```

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
10.207.77.0/24	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
10.1.2.0/24	-	ie0	C	0	0	19775	20887
10.1.2.1/32	-	lo0	CP	0	0	389	20887
255.255.255.255/32	-	ie0	CP	0	0	0	20887

In the output:

- **Destination** is the target address of a route. To send a packet to this address, the Pipeline 220 will use this route. Note that the router will use the most specific route (having the largest subnet mask) that matches a given destination.
- **Gateway** is the address of the next hop router that can forward packets to the given destination. Direct routes (without a gateway) no longer show a gateway address in the gateway column.
- **IF** indicates the name of the interface through which a packet addressed to this destination will be sent.
 ie0 is the Ethernet interface.
 lo0 is the loopback interface.
 wanN specifies each of the active WAN interfaces.
 wanidle0 is the inactive interface (the special interface for any route whose WAN connection is down).
- **Flg** can contain the following flag values:
 - **C** (a directly connected route such as Ethernet)
 - **I** (ICMP Redirect dynamic route)
 - **N** (placed in the table via SNMP MIB II)
 - **O** (a route learned from OSPF)
 - **R** (a route learned from RIP)
 - **r** (a RADIUS route)
 - **S** (a static route)
 - **?** (a route of unknown origin, which indicates an error)
 - **G** (an indirect route via a gateway)
 - **P** (a private route)
 - **T** (a temporary route)
 - ***** (a hidden route that will not be used unless another better route to the same destination goes down)
- **Pref** indicates 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 can be set independently.
- **Metric** shows the RIP-style metric for the route, with a valid range of 0-16. Routes learned from OSPF show a RIP metric of 10. OSPF Cost infinity routes show a RIP metric of 16.
- **Use** is the number of times the route was referenced since it was created. (Many of these references are internal, so this is not a count of the number of packets sent using this route.)
- **Age** is the age of the route in seconds. It is used for troubleshooting, to determine when routes are changing rapidly or flapping.

The first route is the default route (destination 0.0.0.0/0), which is pointing through the active Connection profile:

```
0.0.0.0/0          10.0.0.100      wan0      SG      1      1      0      20887
```

The IP Route profile for the default route specifies a Preference of 1, so this route is preferred over dynamically learned routes. The next route is specified in a Connection profile that is inactive:

```
10.207.76.0/24      10.207.76.1      wanidle0 SG      100      7          0          20887
```

The next route in the table is a static route that points through an inactive gateway:

```
10.207.77.0/24      10.207.76.1      wanidle0 SG      100      8          0          20887
```

The static route is followed by the loopback route:

```
127.0.0.1/32        -                  lo0           CP       0          0          0          20887
```

The loopback route 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.

The next route is specified in a Connection profile that is currently active:

```
10.0.0.0/24          10.0.0.100        wan0          SG       100        1          21387 20887
```

These are followed by the connection to the Ethernet interface. It is directly connected, with a Preference and Metric of zero:

```
10.1.2.0/24          -                  ie0           C        0          0          19775 20887
```

The last two routes are a private loopback route and a private route to the broadcast address:

```
10.1.2.1/32          -                  lo0           CP       0          0          389    20887
255.255.255.255/32 -                  ie0           CP       0          0          0       2088
```

The private loopback route is a host route with the local Ethernet address. It is private, so it will not be advertised. The private route to the broadcast address is used in cases where the router will want to broadcast a packet but is otherwise unconfigured. It is typically used when trying to locate a server on a client machine to handle challenges for a token security card.

Displaying the routing table when using OSPF

The OSPF routing table includes routes built from the router's link-state database, and those added by external routing protocols such as RIP. You can also add routes statically (for example, to direct traffic destined for a remote site through one of several possible border routers). For details about adding static routes, (for example, if you want to force the use of one route over those learned from OSPF), see Chapter 8, "Configuring OSPF Routing."

To view the IP routing table with added OSPF information, enter the `IProute Show` command with the `-l` option:

```
ascend% iproute show -l
```

In addition to the standard routing-table fields, which are described in the *Pipeline 220 Interface Configuration Guide* the following three columns are specific to OSPF and are displayed only when you use the `-l` option. For example, the following OSPF-specific columns are displayed on the far right of each entry in the routing table:

```
...      Cost      T          Tag
...      1         0          0xc0000000
...      9         1          0xc8000000
...     10         0          0xc0000000
```

```

...      9      1      0xc8000000
...      1      1      0xc0000000
...      3      1      0xc8000000
...      9      1      0xc8000000
...      4      1      0xc8000000
...      5      1      0xc8000000
...      3      1      0xc8000000
...      3      1      0xc8000000
...      3      1      0xc8000000

```

In the output:

- **Cost** is the he cost of an OSPF route. The interpretation of this cost depends on type of external metric type, displayed in the next column. If the Pipeline 220 is advertising Type 1 metrics, OSPF can use the specified number as the cost of the route. Type 2 external metrics are an order of magnitude larger.
- **T** is the ASE-type of the metric to be advertised for an external route. Zero (0) in this column indicates an external-type-1 or an OSPF internal route. If this column shows a 1, it means that the route is an external-type-2 route.
- **Tag** is this column specifies a 32-bit hexadecimal number attached to each external route to *tag* it as external to the AS. This number can be used by border routers to filter this record.

Multipath routing

A Pipeline 220 running OSPF can alternate between two equal-cost gateways. When OSPF detects more than one equally good gateway, in terms of routing costs, it puts each equal-cost gateway on an equal-cost list. The router alternates between all the gateways on the list. This is called equal-cost multipath routing. For example, if a router A has two equal-cost routes to *example.com*, one via router B and the other via router C, the following list might result:

Destination	Gateway	IF	Flg	Pref	Met	Use	Age
10.174.88.0/25	10.174.88.12	wan2	OGM	10	10	52	19
10.174.88.0/25	10.174.88.13	wan3	OGM	10	10	52	19
10.174.88.12/32	10.174.88.12	wan2	OG	10	10	0	28
10.174.88.13/32	10.174.88.13	wan3	OG	10	10	0	28
192.168.253.0/24	-	ie0	C	0	0	1	49
192.168.253.6/32	-	lo0	CP	0	0	53	49
223.1.1.0/24	10.174.88.12	wan2	OG	10	10	0	19
223.5.1.0/24	10.174.88.12	wan2	OG	10	10	0	19
223.12.9.0/24	10.174.88.12	wan2	OG	10	10	0	19
255.255.255.255/32	-	ie0	CP	0	0	0	49

M in the Flags column indicates an equal-cost multipath. For example, following is a Traceroute from A to *example.com*:

```

ascend% traceroute -q 10 example.com

traceroute to example.com (10.174.88.1), 30 hops max, 0 byte packets
1  C.example.com (10.174.88.13)  20 ms B .example.com (10.174.88.12)
20 ms C.example.com (10.174.88.13)  20 ms B .example.com
(10.174.88.12)  20 ms  20 ms C.example.com (10.174.88.13)  60 ms  20 ms
B .example.com (10.174.88.12)  20 ms C.example.com (10.174.88.13)  20
ms B .example.com (10.174.88.12)  20 ms

```

```
2  example.com (10.174.88.1)  20 ms  20 ms  20 ms  20 ms  30 ms  20 ms
20 ms  30 ms  20 ms  30 ms
```

Note: Notice the alternating replies. The replies are statistically dispatched to B and C, with roughly 50% of the packets sent through each gateway. For background information about the routing table and about the Traceroute command, see the *Pipeline 220 Interface Configuration Guide*.

Third-party routing

A Pipeline 220 routing OSPF can advertise routes to external destinations on behalf of another gateway (a third-party). This is commonly known as advertising a forwarding address. Depending on the exact topology of the network, it might be possible for other routers to use this type of LSA and route directly to the forwarding address without involving the advertising Pipeline 220, increasing the total network throughput.

Third-party routing requires that all OSPF routers know how to route to the forwarding address. This usually means that the forwarding address must be on an Ethernet that has an OSPF router acting as the forwarding router, or that the designated router is sending LSAs for that Ethernet to any area that sees the static route's forwarding address LSAs.

How OSPF adds RIP routes

When the Pipeline 220 establishes an IP routing connection with a caller that does not support OSPF, it imports the AS-external route from the Connection profile and adds it to the routing table. The Pipeline 220 does not have to run RIP to learn these routes. RIP should be turned off when the Pipeline 220 is running OSPF.

To enable OSPF to add the RIP-v2 routes to its routing table, configure RIP-v2 normally in this Connection profile. OSPF will import all RIP routes as Type-2 ASEs. The reason that RIP routes are imported with Type-2 metrics by default is that RIP metrics are not directly comparable to OSPF metrics. To prevent OSPF from interpreting RIP metrics, the Pipeline 220 assigns the imported ASE route a Type-2 metric, which means that it is so large compared to OSPF costs that the metric can be ignored.

Route preferences

Route preferences provide additional control over which types of routes take precedence over others. They are necessary in a router that speaks multiple routing protocols, largely because RIP metrics are not comparable to OSPF metrics.

For each IP address and subnet mask pair, the routing table holds one route per protocol. 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 OSPF have a default Preference=10.
You can modify the default in Ethernet > Mod Config > Route Pref.
- Routes learned from RIP have a default Preference=100.
You can modify the default in Ethernet > Mod Config > Route Pref.

- A statically configured IP Route or Connection profile has a default Preference=100.
You can modify the default in the Connection or IP Route profile.

When choosing which routes should be put in the routing table, the router first compares the Preference value, preferring the lowest 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 subnet mask pair, the route with the lowest Preference is better. If two routes have the same Preference, then the lower Metric is better. The router uses the best route by these criteria. The others remain latent, or *hidden*, and not used unless the best route is removed.

Adding an IP route

To add to the Pipeline 220's routing table a static route that will be lost when the Pipeline 220 resets, enter the IProute Add command in the following format:

```
iproute add destination gateway [metric]
```

where *destination* is the destination network address, *gateway* is the IP address of the router that can forward packets to that network, and *metric* is the virtual hop count to the destination network (default 8). For example:

```
ascend% iproute add 10.1.2.0 10.0.0.3/24 1
```

The Pipeline 220 adds a route to the 10.1.2.0 network and all of its subnets through the IP router located at 10.0.0.3/24. The metric to the route is 1 (it is one hop away).

If you try to add a route to a destination that already exists in the routing table, the Pipeline 220 replaces the existing route, but only if the existing route has a higher metric. If you get the message `Warning: a better route appears to exist`, the Pipeline 220 rejected your attempt to add a route because the routing table already contained the same route with a lower metric. Note that RIP updates can change the metric for the route.

Deleting an IP route

To remove a route from the Pipeline 220's routing table, enter the IProute Delete command in the following format:

```
iproute delete destination gateway
```

For example:

```
ascend% iproute delete 10.1.2.0 10.0.0.3/24
```

Note: RIP updates can add back any route you remove with IProute Delete. Also, the Pipeline 220 restores all routes listed in the Static Route profile after a system reset.

DNSTab command

The DNSTab command displays DNS-related information. To use the command, include one of two modifiers, which the screen displays when you enter the DNSTab command with a question mark:

```
ascend% dnstab ?  
dnstab ?           Display help information
```

<code>dnstab show</code>	Display local DNS table
<code>dnstab entry</code>	Display local DNS table entry
<code>dnstab edit</code>	Start editor for local DNS table

Monitoring the DNS table

The DNSTab Show command displays the Pipeline 220's DNS table. The command is identical to Show DNSTab. For example:

```
ascend% dnstab show
```

Local DNS Table

Name	IP Address	# Reads	Time of last read
1: " "	-----	-----	
2: "server.corp.com."	200.0.0.0	2	Feb 10 10:40:44
3: "boomerang"	221.0.0.0	2	Feb 10 9:13:33
4: " "	-----	-----	
5: " "	-----	-----	
6: " "	-----	-----	
7: " "	-----	-----	
8: " "	-----	-----	

Displaying specific DNS entries

The DNSTab Entry command displays a list of information about a specific entry in the local DNS table. Enter the DNSTab Entry command in the following format:

```
dnstab entry n
```

where *n* is the number of a DNS-table entry displayed by the DNSTab Show command.

The list includes the entry and all the IP addresses stored for that entry, up to a maximum specified by the Protocols button > IP tab > DNS button > Size parameter.

If you clear the Protocols button > IP tab > DNS button > Enable List Attempt check box, no list appears.

Creating the local DNS table

Use DNSTAB edit to create a local DNS table. The Pipeline 220 disables DNS updates while you use the DNSTab edit command.

The following 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 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.

- 2 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.

- 3 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.
- 4 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.
- 5 When you are finished making entries, type O and press Return when the editor prompts you for another entry.

Editing the local DNS table

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 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.
- 2 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.
- 3 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.
- 4 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.
- 5 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 To display the table, from the terminal server, enter:
`ascend% dnstab edit`
- 2 Type the number of the entry you want to delete and press Return.
- 3 Press the space bar and then press Return.

Note: 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 220 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.

Menu command

The Menu command invokes the terminal-server menu mode, which lists up to four Telnet hosts as configured in Ethernet > Mod Config > TServ Options. Table 16-8 shows a sample listing.

Table 16-8. Sample terminal-server menu

Up to 16 lines of up to 80 characters each will be accepted. Long lines will be truncated. Additional lines will be ignored.
1. host1.abc.com 2. host2.abc.com 3. host3.abc.com 4. host4.abc.com Enter Selection (1-4, q)

To return to the command line, press 0. Terminal-server security must be set up to allow the operator to toggle between the command line and menu mode, or the Menu command has no effect.

Telnet command

The Telnet command initiates a login session to a remote host. Use the following format:

```
telnet [-a|-b|-t] hostname [port-number]
```

If DNS is configured in the Ethernet profile, you can specify a hostname. For example:

```
ascend% telnet myhost
```

If DNS is not configured, you must specify the host's IP address instead. There are also several options in Ethernet > Mod Config > TServer Options that affect Telnet. For example, if Def Telnet is set to Yes, you can just type a hostname to open a Telnet session to that host. For example:

```
ascend% myhost
```

Another way to open a session is to invoke Telnet first, followed by the Open command at the Telnet prompt. For example:

```
ascend% telnet
telnet> open myhost
```

When the Pipeline 220 displays the Telnet prompt, `telnet>`, you can enter any of the Telnet commands described in "Telnet session commands" on page 16-44. You can quit the Telnet session at any time by typing `quit` at the Telnet prompt:

```
telnet> quit
```

Note: During an open Telnet connection, type `Ctrl-]` to display the `telnet>` prompt and the Telnet command-line interface. Any valid Telnet command returns you to the open session. Note that `Ctrl-]` does not function in binary mode Telnet. If you log into the Pipeline 220 by Telnet, you might want to change its escape sequence from `Ctrl-]` to a different setting.

Telnet command arguments

The arguments to the Telnet command are as follows:

Argument	Effect
<i>hostname</i>	If DNS is configured, you can specify the remote system's hostname. Otherwise, hostname must be the IP address of the remote station.
<code>-a -b -t</code>	(Optional.) You can specify <code>-a</code> , <code>-b</code> , or <code>-t</code> on the Telnet command line to indicate ASCII, Binary, or Transparent mode, respectively. A specification on the command line overrides the setting of the Telnet Mode parameter. The modes have the following effects: <ul style="list-style-type: none">• In ASCII mode, the Pipeline 220 uses standard 7-bit mode.• In Binary mode, the Pipeline 220 tries to negotiate 8-bit Binary mode with the server at the remote end of the connection.• In Transparent mode, the user can send and receive binary files, and use 8-bit file transfer protocols, without having to be in Binary mode.
<i>port-number</i>	(Optional.) You can specify the port to use for the session. The default is 23, the well-known port for Telnet.

Telnet session commands

The commands in the following section can be typed at the Telnet prompt during an open session. To display the Telnet prompt during an active login to the specified host, press `Ctrl-]`

(hold down the Control key and type a right-bracket). To display information about Telnet session commands, use the Help or ? command:

```
telnet> ?
```

To open a Telnet connection after invoking Telnet, use the Open command; for example:

```
telnet> open myhost
```

To send standard Telnet commands such as Are You There or Suspend Process, use the Send command. For example:

```
telnet> send susp
```

For a list of Send commands and their syntax, type:

```
telnet> send ?
```

To set special characters for use during the Telnet session, use the SET command. For example:

```
telnet> set eof ^D
```

To display current settings, enter the Set All command:

```
telnet> set all
```

To view a list of Set commands, enter the Set command with a question mark:

```
telnet> set ?
```

To quit the Telnet session and close the connection, use the Close or Quit command. For example:

```
telnet> close
```

Telnet error messages

The Pipeline 220 generates an error message for any condition that causes the Telnet session to fail or terminate abnormally. The following error messages might appear:

- no connection: host reset (The destination host reset the connection.)
- no connection: host unreachable (The destination host is unreachable.)
- no connection: net unreachable (The destination network is unreachable.)
- Unit busy. Try again later. (The host already has the maximum number of concurrent Telnet sessions open.)

TCP command

The TCP command initiates a login session to a remote host. Use the following format:

```
tcp hostname port-number
```

where:

- *hostname* is the remote system's hostname. You can specify it if DNS is configured. Otherwise, hostname must be the IP address of the remote station.
- *port-number* (Optional.) specifies the port to use for the session. The port number typically indicates a custom application that runs on top of the TCP session. For example,

port number 23 starts a Telnet session. However, terminating the Telnet session does not terminate the raw TCP session.

For example:

```
ascend% tcp myhost
```

When the raw TCP session starts running, the Pipeline 220 displays the word *connected*. You can now use the TCP session to transport data by running an application on top of TCP. You can hang up the device at either end to terminate the raw TCP session. If you are using a remote terminal-server session, ending the connection also terminates raw TCP.

If a raw TCP connection fails, the Pipeline 220 returns one of the following error messages:

Can't open session: *hostname port-number*

You entered an invalid or unknown value for *hostname*, you entered an invalid value for *port-number*, or you failed to enter a port number.

- no connection: host reset (The destination host reset the connection.)
- no connection: host unreachable (The destination host is unreachable.)
- no connection: net unreachable (The destination network is unreachable.)

Ping command

The terminal-server Ping command is useful for verifying that the transmission path is open between the Pipeline 220 and another station. It sends an ICMP echo_request packet to the specified station. If the station receives the packet, it returns an ICMP echo_response packet. For example, to ping the host techpubs:

```
ascend% ping techpubs
PING techpubs (10.65.212.19): 56 data bytes
64 bytes from 10.65.212.19: icmp_seq=0 ttl=255 time=0 ms
64 bytes from 10.65.212.19: icmp_seq=3 ttl=255 time=0 ms
^C
--- techpubs ping statistics ---
2 packets transmitted, 2 packets received, 0% packet loss
round-trip min/avg/max = 0/0/0 ms
```

You can terminate the Ping exchange at any time by typing Ctrl-C. When you press Ctrl-C, the command reports the number of packets sent and received, the percentage of packet loss, duplicate or damaged echo_response packets (if any), and round-trip statistics. In some cases, round-trip times cannot be calculated.

During the Ping exchange, the Pipeline 220 displays information about the packet exchange, including the TTL (Time-To-Live) of each ICMP echo_response packet.

Note: The maximum TTL for ICMP Ping is 255, and the maximum TTL for TCP is often 60 or lower, so you might be able to ping a host but not be able to run a TCP application (such as Telnet or FTP) to that station. If you Ping a host running a version of Berkeley UNIX earlier than 4.3BSD-Tahoe, the TTL report is 255 minus the number of routers in the round-trip path. If you Ping a host running the current version of Berkeley UNIX, the TTL report is 255 minus the number of routers in the path from the remote system to the station performing the Ping.

The Ping command sends an 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 220 and a remote station.

IPXping command

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 220 or across a WAN connection that has IPX Routing enabled. Use the following format:

```
ipxping [-c count] [-i delay] [-s packetsize] hostname
```

The arguments to the IPXping command are:

- *hostname*: The IPX address of the host, or if the host is a NetWare server, its hostname.
- *[-c count]*: (Optional): Stop the test after sending and receiving the number of packets specified by count.
- *[-i delay]*(Optional): Wait the number of seconds specified by wait before sending the next packet. The default is one second.
- *[-s packetsize]*(Optional): Send the number of data bytes specified by packet-size.

The *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. For example:

```
ascend% ipxping CFFF1234:000000000001
```

If you are using IPXping to verify connectivity with an advertised NetWare server, you can simply enter the symbolic name of the server. For example:

```
ascend% ipxping server-1
```

You can terminate the IPXping at any time by typing Ctrl-C.

During the IPXping exchange, the Pipeline 220 calculates and reports the following information:

```
PING server-1 (EE000001:000000000001): 12 data bytes
52 bytes from (EE000001:000000000001): ping_id=0 time=0ms
52 bytes from (EE000001:000000000001): ping_id=1 time=0ms
52 bytes from (EE000001:000000000001): ping_id=2 time=0ms
?
--- nov11 Ping statistics ---
3 packets transmitted, 3 packets received, 0% packet loss
round-trip min/avg/max = 0/0/0 ms
```

The output includes the following information:

- The IPX address of the source and destination nodes.
- The byte counts of the request and response packets.
- The ping ID of the command. (The ping Request # replied to by target host.)
- The number of milliseconds required to send the IPXping and receive a response.
- The number of packets transmitted and received.
- Duplicate or damaged packets, if applicable.

- Average round-trip times for the ping request and reply.

In some cases, round-trip times cannot be calculated.

To display statistics related to the IPXping command, type:

```
ascend% show network pings
InPing Requests/OutPing Replies OutPing Requests/InPing Replies
      10              10              18              18
```

The output shows how many NetWare stations have pinged the Pipeline 220 (InPing requests and replies) and how many times the IPXping command has been executed in the Pipeline 220.

Traceroute command

The Traceroute command is useful for locating slow routers or diagnosing IP routing problems. It traces the route an IP packet follows by launching UDP probe packets with a low TTL (Time-To-Live) value and then listening for an ICMP time exceeded reply from a router. The format of the Traceroute command is as follows:

```
traceroute [ -n ] [ -v ] [ -m max_ttl ] [ -p port ] [ -q nqueries ]
[ -w waittime ] host [ datasize ]
```

All flags are optional. The only required parameter is the destination hostname or IP address. The flags have the following effects:

Flag	Effect
-n	Prints hop addresses numerically rather than symbolically and numerically (this eliminates a name server address-to-name lookup for each gateway found on the path).
-v	Verbose output. Received ICMP packets other than Time Exceeded and ICMP Port Unreachable are listed.
-m <max_ttl>	This sets the maximum time-to-live (maximum number of hops) used in outgoing probe packets. The default is 30 hops.
-p <port>	Sets the base UDP port number used in probes. Traceroute hopes that nothing is listening on any of the UDP ports from the source to the destination host (so an ICMP Port Unreachable message will be returned to terminate the route tracing). If something is listening on a port in the default range, this option can be used to pick an unused port range. The default is 33434.
-q <nqueries>	Sets the maximum number of queries for each hop. The default is 3.
-w <waittime>	Sets the time to wait for a response to a query. The default is 3 seconds.
host	The destination host by name or IP address.
datasize	Sets the size of the data field of the UDP probe datagram sent by Traceroute. The default is 0. This results in a datagram size of 38 bytes (a UDP packet carrying no data).

For example, to trace the route to the host techpubs:

```
ascend% traceroute techpubs
traceroute to techpubs (10.65.212.19), 30 hops max, 0 byte packets
 1  techpubs.eng.ascend.com (10.65.212.19)  0 ms  0 ms  0 ms
```

Probes start with a TTL of one and increase by one until one of the following conditions occurs:

- The Pipeline 220Pipeline 220 receives an ICMP `port unreachable` message.
The UDP port in the probe packets is set to an unlikely value, such as 33434, because the target host is not intended to process the packets. A `port unreachable` message indicates that the packets reached the target host and were rejected.
- The TTL value reaches the maximum value.
By default, the maximum TTL is set to 30. You can specify a different TTL by using the `-m` option. For example:

```
ascend% traceroute -m 60 techpubs
```

```
traceroute to techpubs (10.65.212.19), 60 hops max, 0 byte packets
 1 techpubs.eng.abc.com (10.65.212.19) 0 ms 0 ms 0 ms
```

Three probes are sent at each TTL setting. The second line of command output shows the address of the router and round trip time of each probe. If the probe answers come from different gateways, the address of each responding system will be printed. If there is no response within a 3 second timeout interval, the command output is an asterisk. The following annotations might be included after the time field in a response:

- `!H` (Host reached.)
- `!N` (Network unreachable.)
- `!P` (Protocol unreachable.)
- `!S` (Source route failed. This might indicate a problem with the associated device.)
- `!F` (Fragmentation needed. This might indicate a problem with the associated device.)
- `!h` (Communication with the host is prohibited by filtering.)
- `!n` (Communication with the network is prohibited by filtering.)
- `!c` (Communication is otherwise prohibited by filtering.)
- `!?` (Indicates an ICMP sub-code. This should not occur.)
- `!??` (Reply received with inappropriate type. This should not occur.)

Kill command

The Kill command enables you to clear the nailed connection. Disconnect by using the session ID. The disconnect code that results is identical to the RADIUS disconnect code, allowing you to track all administrative disconnects. Use the following format:

```
kill session ID
```

where *session ID* is the session ID as displayed by the Show Users command described in the preceding section. The reported disconnect cause is `DIS_LOCAL_ADMIN`. The active Security profile must have Edit All Calls=Yes. If Edit All Calls=No, the Pipeline 220 displays the following message when you issue the kill command:

```
Insufficient security level for that operation.
```

When the session is properly terminated, the Pipeline 220Pipeline 220 displays the following message:

```
Session 216747095 killed.
```


Pipeline 220 Specifications

A

This appendix describes specifications for different facets of the Pipeline 220, and discusses cabling requirements. The appendix contains the following sections:

General specifications	A-1
User interface specifications	A-3
Ethernet interface specifications	A-3
Serial WAN cabling specifications	A-4

General specifications

Battery

The Pipeline 220 contains an internal 3V lithium battery. The normal operating life of this battery exceeds five years.

Only trained engineers authorized by Ascend should open the Pipeline 220's case for testing, maintenance, installation, or any other purpose. Furthermore, only trained engineers should replace Pipeline 220 components.



Warning: The battery can explode if incorrectly replaced. Replace the battery only with the same or equivalent type recommended by the manufacturer. Dispose of used batteries according to the manufacturer's instructions.

ATTENTION: IL Y A DANGER D'EXPLOSION S'IL Y A REMPLACEMENT INCORRECT DE LA BATTERIE. REMPLACER UNIQUEMENT AVEC UNE BATTERIE DU MÊME TYPE OU D'UN TYPE RECOMMANDÉ PAR LE CONSTRUCTEUR. METTRE AU RÉBUT LES BATTERIES USAGÉES CONFORMÉMENT AUX INSTRUCTIONS DU FABRICANT.

Power requirements

Table A-1 lists the Pipeline 220's source power requirements.

Table A-1. Pipeline 220 source power requirements

Element	Value
Voltage	90-240 VAC
Phase	Single
Frequency	47-63 Hz
Power	80W (nominal)-200W (maximum)

The Pipeline 220's configuration profiles are stored in battery-protected memory. When the Pipeline 220 is turned off, the profiles are not lost.

Note: Use a protected AC power source, or add surge protection between the power source and the Pipeline 220.

Environmental requirements

For best results, you should house the Pipeline 220 in a room with constant temperature and humidity. In general, cooler environments are better, and an operating temperature of 32° to 104° Fahrenheit (0° to 40° Celsius) is recommended. Storage temperatures of -40° to 176° Fahrenheit (-71.4° to 80° Celsius) are acceptable.

Humidity should be high enough to prevent accumulation of static electricity, but low enough to prevent condensation. An operating relative humidity of up to 90% is acceptable.

You can operate the Pipeline 220 at altitudes of from 0 to 14800 ft. (0-4500 m).

The Pipeline 220 base system weighs 7.25 lbs (3.3 kg). The Pipeline 220 has dimensions of: 17.63 in x 2.0 in x 8.25 in (44.8 cm x 5.1 cm x 21.0 cm).

User interface specifications

This sections covers cabling pinouts for the Control Monitor. The Control port uses a standard DE-9 female connector that conforms to the EIA RS-232 standard for serial interfaces. All Pipeline 220 models use the RS-232 pinouts listed in Table A-2..

Table A-2. Control Monitor port cabling pinouts

DE-9 pin number	RS-232 signal name	Function	I/O
1	DCD	Data Carrier Detect	O
2	RD	Serial Receive Data	O
3	SD	Serial Transmit Data	I
4	DTR	Data Terminal Ready	I
5	GND	Signal Ground	
6	DSR	Data Set Ready	O
7	RTS	Request to Send	I
8	CTS	Clear to Send	O
*9	*RI	*Ring Indicator	*O

*Pin 9 is not active (Ring Indication signal not supplied).

Ethernet interface specifications

The two Ethernet interfaces of the Pipeline 220 support the physical specifications of IEEE 1802.3 with Ethernet 2 (Ethernet/DIX) framing. Each Ethernet port provides a single Ethernet interface which supports 10Base-T (Unshielded Twisted Pair), which is also referred to as Twisted pair Ethernet and IEEE 802.3 (10BaseT) with an RJ-45 connector.

Serial WAN cabling specifications

The Pipeline 220's serial WAN interface supports nailed-up connections to the WAN. Data packets from the Pipeline 220's bridge/router module can use this interface.

The Pipeline 220's serial WAN port is compatible with the following two electrical standards:

- V.35
- RS-449/422

In the cable wiring tables that follow, the Pipeline 220 is the Data Terminal Equipment (DTE) that connects to a Data Communications Equipment (DCE) device through its serial WAN port. The Pipeline 220 receives the Send timing and Receive timing clocks from the DCE device.

V.35 cable to WAN

You connect a V.35 cable to the V.35 port of a DTE device. If you order the cable, you also received a female-to-male V.35 gender changer. The V.35 cable has the pinouts shown in Table A-3.

Table A-3. V.35 cable pinouts

Part number 2510-0202-001			
Pair #	Signal	Pipeline 220 male DB-44	Host male V.35
1	FGND	1	A
	RI	8	J
2	SD+	39	P
	SD-	40	S
3	RD+	30	R
	RD-	29	T
4	ST+	41	Y
	ST-	42	AA
5	RT+	32	V
	RT-	31	X
6	TT+	38	U
	TT-	37	W
7	DTR	6	H
	DSR	11	E

Table A-3. V.35 cable pinouts (continued)

Part number 2510-0202-001			
Pair #	Signal	Pipeline 220 male DB-44	Host male V.35
8	DCD	9	F
	SGND	25	B
9	CTS	7	D
	RTS	36	C

RS-449 cable to WAN

If you order the cable, you also receive a female-to-male DB-37 gender changer. The RS-449 cable has the pinouts shown in Table A-4.

Table A-4. RS-449 cable pinouts

Part number 2510-0203-001			
Pair #	Signal	Pipeline 220 male DB-44	Host female DB-37
1	FGND	1	1
	RI	8	15
2	SD+	39	4
	SD-	40	22
3	RD+	30	6
	RD-	29	24
4	ST+	41	5
	ST-	42	23
5	RT+	32	8
	RT-	31	26
9	TT+	38	17
	TT-	37	35
8	DTR	6	12
	DSR	11	11
6	DCD	9	13
	SGND	25	19, 20, 37*

Table A-4. RS-449 cable pinouts (continued)

Part number 2510-0203-001			
Pair #	Signal	Pipeline 220 male DB-44	Host female DB-37
7	CTS	7	9
	RTS	36	7

* Pin positions separated by commas are jumped to each other.

Upgrading System Software

B



Caution: You must use the new software loading procedure explained in "Upgrading system software" to load this version of software onto your system. Read the instructions carefully before upgrading your system.

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.

Each incremental release contains new features and corrections. To use this release note:

- 1 Read through the table of contents to determine which software release and (new features) apply to your environment.
- 2 Obtain the file from Ascend anonymous FTP server (<ftp.ascend.com>). If you need Technical Assistance, 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)
- UK	(+33) 492 96 5671
- Germany/Austria/Switzerland	(+33) 492 96 5672
- France	(+33) 492 96 5673
- Benelux	(+33) 492 96 5674
- Spain/Portugal	(+33) 492 96 5675
- Italy	(+33) 492 96 5676
- Scandinavia	(+33) 492 96 5677
- Middle East and Africa	(+33) 492 96 5679
E-mail	support@ascend.com
E-mail (outside US)	EMEAsupport@ascend.com
Facsimile (FAX)	510-814-2312
Customer Support BBS by modem	510-814-2302

- 3 Upgrade to the new software by following the instructions in the next section, "Upgrading system software." Then configure the features that apply to your site.

Upgrading system software



Caution: The procedure for uploading new software to Ascend units have changed significantly. Carefully read the new software loading procedures explained in this section before upgrading your system.

This section explains how to upgrade your system software. It contains the following sections:

- Definitions and terms
- Guidelines for upgrading system software
- Before you begin
- Upgrading system software
- Using TFTP to upgrade to a fat or thin load
- Upgrading software with an extended load
- Upgrading software from versions earlier than 4.6C to version 5.0A or above
- Using the serial port to upgrade to a standard or a thin load
- System messages

Definitions and terms

This document uses the following terms:

Build	<p>The name of the software binary.</p> <p>For example, <code>ti.m40</code> is the MAX 4000 T1 IP-only software build. For the names of all the software builds and the features they provide see</p> <p><code>/pub/Software-Releases/Max/SW-Filenames-Max.txt</code> or</p> <p><code>/pub/Software-Releases/Pipeline/SW-Filenames-Pipeline.txt</code> on the Ascend FTP server.</p> <p>If possible, you should stay with the same build when upgrading. Loading a different build can cause your Ascend unit to lose its configuration. If this happens, you must restore your configuration from a backup.</p>
Standard load	<p>Software versions 4.6Ci18 or earlier and all 4.6Cp releases. You can load these versions of software through the serial port or by using TFTP. TFTP is the recommended upgrade method for standard loads.</p>
Fat load	<p>4.6Ci19 to 5.0Aix and all 5.0Ap releases with a file size greater than 960 KB (for MAX units) or 450K (for Pipeline units). Before upgrading to a fat load for the first time, you must upgrade to a thin load. You must use TFTP to upgrade to fat loads.</p>
Thin load	<p>4.6Ci19 to 5.0Aix and all 5.0Ap releases with a file size less than 960 KB (for MAX units) or 450 KB (for Pipeline units). TFTP is the recommended upgrade method for thin loads.</p>

Restricted load	6.0.0 or later release denoted by an “r” preceding the build name. For example, <code>rti.m40</code> is the restricted load for the MAX 4000 T1 IP-only software build A restricted load is not meant for production environments. It is a special load that is required to upgrade to an extended load. Before upgrading to an extended load for the first time, you must upgrade to a restricted load. You must use TFTP to upgrade to restricted loads.
Extended load	6.0.0 or later release. You must use TFTP to upgrade to extended loads.

Guidelines for upgrading system software



Caution: Before upgrading, consider the following very important guidelines:

- Use TFTP to upgrade if possible. TFTP is more reliable and saves the Ascend unit configuration when you upgrade.
- You cannot load or a fat load, a restricted load, or an extended load through the serial port. You must use TFTP.
- If you are using TFTP to upgrade your software, use the `fsave` command immediately after executing the `tload` command. Failure to do so might cause your Ascend unit to lose its configuration.
- If possible, you should always stay with the same build of software when you upgrade. If you load a different version, your Ascend unit may lose its configuration. If this happens, you must restore your configuration from a backup.
- If you are upgrading to a software version 5.0A or 5.0Aix fat load for the first time, you must be on a load that supports the fat load format. All versions of software 5.0A or above support fat loads. You should perform the upgrade in two steps:
 - Upgrade to a thin load of the same build
 - Upgrade to the fat load
- If you are upgrading to a software version 6.0.0 or above, you must be on a load that supports the extended load format. All versions of software 6.0.0 or above support extended loads. You should perform the upgrade in two steps:
 - Upgrade to a restricted load of the same build
 - Upgrade to the extended load
- You can upgrade to a thin load or a restricted load from any version of software.
- If you are upgrading from software version 4.6C or earlier to software version 5.0A or later, see “Upgrading software from versions earlier than 4.6C to version 5.0A or above” on page B-10 for important information before you start.

Table B-1 explains where to find the information you need to upgrade your unit.

Table B-1. Ascend system software versions

Version you are upgrading to	Use the instructions in...
Standard load (4.6Ci18 or earlier and all 4.6Cp releases)	“Upgrading system software” on page B-5.
Fat or thin load (4.6Ci19 to 5.0Aix and all 5.0Ap releases)	“Using TFTP to upgrade to a fat or thin load” on page B-6.
Extended load (6.0.0 or later)	A restricted load is not meant for production environments. It is a special load that is required to upgrade to an extended load. “Upgrading software with an extended load” on page B-9.

Before you begin

Make sure you perform all the tasks explained in Table B-2 before upgrading your software.

Table B-2. Before upgrading

Task	Description
If necessary, activate a Security Profile that allows for field upgrade.	If you are not sure how, see the section about Security Profiles in your documentation.
Record all of the passwords you want to retain, and save your Ascend unit’s current configuration to your computer’s hard disk.	For security reasons, passwords are not written to configuration files created through the serial console. A configuration file created using the Tsave command, however, <i>does</i> contain the system passwords. You can restore the Tsave configuration file using the serial console. If you chose to save your configuration using the serial console, you will have to restore your passwords manually. Restoring passwords is explained in “Using the serial port to upgrade to a standard or a thin load” on page B-11.

Upgrading system software

Table B-2. Before upgrading (Continued)

Task	Description
Obtain the correct file, either by downloading it from the FTP server or by requesting it from Ascend technical support.	<p>To ensure that you load the correct software binary, you should check the load currently installed on your unit. To do so:</p> <ol style="list-style-type: none"> 1 Tab over to the System status window. 2 Press Enter to open the Sys Options menu. 3 Using the Down-Arrow key (or Ctrl- N), scroll down until you see a line similar to the following: Load: tb.m40 4 When upgrading, obtain the file with same name from the Ascend FTP site. <p>If your unit does not display the current load or you are unsure about which load to use, contact technical support.</p>
If you are upgrading to a fat load or an extended load for the first time, you must also obtain a thin load or a restricted load of the same build.	<p>For example, if you are upgrading a MAX 4000 to 5.0Ai13 fat load (such as tbim.m40), obtain a thin load of the same build (such as 5.0A tbim.m40).</p> <p>If you are upgrading to a 6.0.0 extended load, obtain a 6.0.0 restricted load. Restricted loads are designated with an “r” in the load name. (For example rtbam.m40 is a restricted load).</p> <p>Newer Pipeline 50 or 75 units do not have fat or extended loads. Refer to the README file in /pub/Software-Releases/Pipeline/software-version to determine if you have a new Pipeline 50 or 75 unit.</p>
If you are using TFTP, make sure you load the correct binaries into the TFTP home directory on the TFTP server.	You must use TFTP to upgrade to a fat load or an extended load.
If you are using the serial port, make sure you have a reliable terminal emulation program, such as Procomm Plus.	<p>If you use the serial port, you can only upgrade to a standard or a thin load. Upgrading through the serial port is not recommended.</p> <p>If you use a Windows-based terminal emulator such as Windows Terminal or HyperTerminal, disable any screen savers or other programs or applications that could interrupt the file transfer. Failure to do so might cause the software upload to halt, and can render the Ascend unit unusable.</p>

Upgrading system software

To upgrade system software with a standard load you can use either the serial port or TFTP. TFTP is the recommended method because it preserves your Ascend unit’s configuration. If you want to use the serial port to upgrade, see “Using the serial port to upgrade to a standard or a thin load” on page B-11.

Using TFTP to upgrade to a standard load

To upgrade to a standard load using TFTP, you only have to enter a few commands. But you must enter them in the correct sequence, or you could lose the Ascend unit's configuration.

To upgrade to a standard load via TFTP:

- 1 Obtain the software version you want to upgrade to and place it in the TFTP server home directory.
- 2 From the Ascend unit's VT100 interface, access the diagnostics monitor by typing the following characters in rapid succession:

`ESC [ESC =`

Or, press Ctrl-D to invoke the DO menu and select D=Diagnostics.

- 3 At the > prompt, use the Tsave command to save your configuration as in the following example:

> tsave tftp-server router1.cfg

This saves the configuration of your unit to the file named `router1.cfg` in the TFTP home directory of the server named `tftp-server`. This file must already exist and be writable. Normally, TFTP upgrades save the configuration. Tsave is a precaution.



Caution: The file you save with the Tsave command contains all the passwords in clear text. You should move this file from the TFTP directory to a secure location after the upgrade procedure is complete.

- 4 Enter the following command:

tloadcode *hostname filename*

where *hostname* is the name or IP address of your TFTP server, and *filename* is the name of the system software on the server (relative to the TFTP home directory).

For example, the command:

tloadcode tftp-server t.m40

loads `t.m40` into flash from the machine named `tftp-server`.



Caution: You must use the Fsave command immediately after executing the Tload command. Failure to do so can cause your Ascend unit to lose its configuration.

- 5 Enter the following command to save your configuration to flash memory:

fsave

- 6 Enter the following command:

nvrnrclear

After the Ascend unit clears NVRAM memory, it automatically resets.

This completes the upgrade.

Using TFTP to upgrade to a fat or thin load

Upgrading to a fat or thin load is not difficult, but you must be careful to follow the correct sequence of tasks.



Caution: If you are upgrading from software version 4.6C or earlier, see “Upgrading software from versions earlier than 4.6C to version 5.0A or above” on page B-10 for important information before upgrading.

To upgrade your system:

- 1 Obtain the software version binary you want to upgrade to and place it in the TFTP server home directory. If you are upgrading to a fat load for the first time, also obtain a thin load of the same build and place it in the same directory.



Caution: If possible, you should stay with the same build when upgrading. Loading a different build can cause your Ascend unit to lose its configuration. If this happens, you must restore your configuration from a backup.

For example, if you are upgrading a MAX 4000 to 5.0Ai13 fat load (such as `tbim.m40`), obtain a thin load of the same build (such as `5.0A tbim.m40`).

Note: Newer Pipeline 50 or 75 units do not have fat or thin loads, you only need to load a single software binary. Refer to the README file in `/pub/Software-Releases/Pipeline/software-version` on the Ascend FTP site to determine if you have a new Pipeline 50 or 75 unit.

- 2 From the Ascend unit's VT100 interface, access the diagnostics monitor by typing the following characters in rapid succession:
`Esc [Esc =`
 Or, press Ctrl-D to invoke the DO menu and select D=Diagnostics.
- 3 At the `>` prompt, use the `Tsave` command to save your configuration, as in the following example:

```
> tsave tftp-server router1.cfg
```

This saves the configuration of your unit to the file named `router1.cfg` in the TFTP home directory of the server named `tftp-server`. This file must already exist and be writable. Normally, TFTP upgrades save the configuration. `Tsave` is a precaution.



Caution: The file you save with the `Tsave` command contains all the passwords in clear text. You should move this file from the TFTP directory to a secure location after the upgrade procedure is complete.

- 4 At the `>` prompt, enter:
`tloadcode hostname filename`
 where *hostname* is the name or IP address of your TFTP server, and *filename* is the name of the system software on the server (relative to the TFTP home directory).



Caution: If you are upgrading from a standard load to a fat load, make sure you load a thin load first.

For example, the command:

```
tloadcode tftp-server t.m40
```

loads `t.m40` into flash from the machine named `tftp-server`.



Caution: 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.

- 5 Enter the following command to save your configuration to flash memory:
fsave
- 6 Enter the following command:
nvrampclear
After the Ascend unit clears NVRAM memory, it automatically resets.
- 7 If you are upgrading to a thin load, you are done. If you are upgrading to a fat load, repeat the procedure, this time uploading the fat load binary.

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...

This completes the upgrade if you have no errors. If the upgrade is not successful, refer to "Recovering from a failed fat load upgrade" next.

Recovering from a failed fat load upgrade

If a fat load has a CRC (cyclic redundancy check) 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 Ascend unit initiates an Xmodem serial download. To recover from this error and load the fat system, you must first load a thin system that is fat-load aware. Proceed as follows:

- 1 Activate your Xmodem software.
- 2 After you have finished loading the fat-aware thin load, reboot the unit.
- 3 Use the Tload command to download the fat load.
When you download a fat load, messages similar to the following appear on the diagnostics monitor screen:

```
> tload 192.168.1.82 tbam.m40
saving config to flash
.....
loading code from 192.168.1.82:69
file tbam.m40..
fat load part 1:
.....
```

```
.....
fat load part 2:
.....
```

The “fat load part *n*:” messages notify you when the first and second halves of the download begin.

Upgrading software with an extended load

Your first upgrade to an extended load requires a preliminary procedure. You must first upgrade to a restricted load. Restricted loads are not meant to be used in a working unit. They are a temporary load that are only used to prepare your Ascend unit for the extended load.



Caution: If you are upgrading from software version 4.6C or earlier, see “Upgrading software from versions earlier than 4.6C to version 5.0A or above” on page B-10 for important information before upgrading.

To upgrade your system:

- 1 Obtain the software-version binary you want to upgrade to and place it in the TFTP server home directory.
- 2 If this is the first time you have upgraded to an extended load, obtain a restricted load of the same build and place it in the directory.
For example, if you are upgrading a MAX 4000 to an extended load (such as `tbam.m40`), obtain a MAX 4000 restricted load (such as `rtbam.m40`).

Note: Newer Pipeline 50 or 75 units do not have restricted loads, you only need to load a single software binary. Refer to the README file in `/pub/Software-Releases/Pipeline/software-version` on the Ascend FTP site to determine if you have a new Pipeline 50 or 75 unit.

- 3 From the Ascend unit’s VT100 interface, access the diagnostics monitor by typing the following characters in rapid succession:
`Esc [Esc =`
Or, press Ctrl-D to invoke the DO menu, and select D=Diagnostics.
- 4 At the `>` prompt, use the `Tsave` command to save your configuration, as in the following example:

```
> tsave tftp-server router1.cfg
```

This saves the configuration of your unit to the file named `router1.cfg` in the TFTP home directory of the server named `tftp-server`. This file must already exist and be writable. Normally, TFTP upgrades save the configuration. `Tsave` is a precaution.



Caution: The file you save with the `Tsave` command contains all the passwords in clear text. You should move this file from the TFTP directory to a secure location after the upgrade procedure is complete.

- 5 At the `>` prompt, enter:
`tloadcode hostname filename`
where *hostname* is the name or IP address of your TFTP server, and *filename* is the name of the system software on the server (relative to the TFTP home directory).



Caution: If you want to upgrade your system for the first time to a software version 6.0.0 or later, you must first upgrade your system to a restricted load. Failure to do so can cause your Ascend unit to lose its configuration.

For example, the command:

```
tloadcode tftp-server rtbam.m40
```

loads the restricted load `rtbam.m40` into flash from the machine named `tftp-server`.



Caution: You must use the `Fsave` command immediately after executing the `Tload` command. Failure to do so can cause your Ascend unit to lose its configuration.

- 6 Enter the following command to save your configuration to flash memory:

```
fsave
```

- 7 Enter the following command:

```
nvrampclear
```

After the Ascend unit clears NVRAM memory, it automatically resets.

If you have downloaded the extended load, the upgrade is complete.

If you have loaded a restricted load, your system boots up in restricted mode. Restricted mode only allows you to load software. You cannot change or save profiles. While in restricted mode, the Edit menu displays the following banner:

```
* * RESTRICTED MODE * * * YOU MUST RERUN THE LAST tloadcode COMMAND *  
*
```

If your system boots up in restricted mode, repeat the step 5 through step 7 to download the extended load.

Upgrading software from versions earlier than 4.6C to version 5.0A or above

If you are upgrading from software version 4.6C or earlier to version 5.0A or later, perform the upgrade in the following order:

- 1 Load version 4.6Ci18, following the procedure in “Upgrading system software” on page B-5.
- 2 Load version 5.0A, following the procedure in “Using TFTP to upgrade to a fat or thin load” on page B-6.
- 3 Load version 5.0Aix or 6.0.0, following the procedure in “Using TFTP to upgrade to a fat or thin load” on page B-6 (for software versions 5.0Aix) or “Upgrading software with an extended load” on page B-9 (for software version 6.0.0).



Caution: Failure to follow this procedure might cause your Ascend unit to lose or corrupt its configuration, and could render the unit unusable.

Using the serial port to upgrade to a standard or a thin load



Caution: Uploading system software via the serial console overwrites all existing profiles. Save your current profiles settings to your hard disk before you begin upgrading system software. After the upgrade, restore your profiles from the backup file you created. Since the backup file is readable text, you can reenter the settings through the Ascend unit's user interface. To avoid having existing profiles overwritten, use TFTP to upgrade your unit.



Caution: You cannot upload a fat load or an extended load using the serial port; it must be done using TFTP.

Upgrading through the serial port consists of the following general steps:

- Saving your configuration
- Uploading the software
- Restoring the configuration

Before you begin

Before upgrading your system through the serial port, make sure you have the following equipment and software:

- An IBM compatible PC or Macintosh with a serial port capable of connecting to the Ascend unit's Console port.
- A straight-through serial cable.
- Data communications software for your PC or Mac with XModem CRC/1K support (for example, Procomm Plus, HyperTerminal for PCs or ZTerm for the Mac).



Caution: If you use a Windows-based terminal emulator such as Windows Terminal or HyperTerminal, disable any screen savers or other programs or applications that could interrupt the file transfer. Failure to do so might cause the software upload to halt, and can render the Ascend unit unusable.

Saving your configuration

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 the same rate as the Term Rate parameter in the System Profile (Sys Config menu).

You can cancel the backup process at any time by pressing Ctrl-C.

To save the Pipeline configuration (except passwords) to disk:

- 1 Open the Sys Diag menu.
- 2 Select Save Config, 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.)
- 4 Press any key to start saving your configured profiles.
Rows of configuration information appear on the screen as the configuration file is downloaded to your hard disk. When the file has been saved, 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.

You should examine the saved configuration file. Notice that some of the lines begin with `START=` and other lines begin with `END=`. A pair of these `START/STOP` lines and the block of data 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, in which case the corresponding `START/STOP` blocks are 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 Ascend unit.

Uploading the software

To upload the software:

- 1 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 appears:
`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.
- 2 Use the Xmodem file-transfer protocol to send the system file to the Ascend unit.
Your communications program normally takes anywhere from 5 to 15 minutes to send the file to your Ascend unit. 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.

After the upload, the Ascend unit resets. Upon completion of the self-test, the Ascend unit’s initial menu appears in the Edit window with all parameters set to default values. This completes the upgrade.

If the upload fails during the transfer, try downloading another copy of the binary image from the Ascend FTP server and re-loading the code to the Ascend unit. If you still have problems, contact Ascend technical support for assistance.

Restoring the configuration

Under certain circumstances, the serial-port method might not completely restore your configuration. You should therefore verify that your configuration was properly restored every time you use this method. If you have many profiles and passwords, you should consider using

TFTP to upgrade your software. (See “Using TFTP to upgrade to a standard load” on page B-6.)

To restore the configuration, you must have administrative privileges that include Field Service (such as the Full Access Profile, for example). You 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 through the serial port

- 1 From the Ascend unit’s VT100 interface, access the diagnostics monitor by typing the following characters in rapid succession:

`ESC [ESC =`

Or, press Ctrl-D to invoke the DO menu, and select D=Diagnostics.

- 2 At the > prompt, enter the Fclear command:

`> fclear`

- 3 At the > prompt, enter the NVRAMClear command:

`> nvramclear`

This causes the system to reset. When it comes back up, proceed with restoring your configuration.

- 4 Enter **quit** to exit the Diagnostic interface.

- 5 Open the Sys Diag menu.

- 6 Select Restore Cfg, and press Enter.

The following message appears:

`Waiting for upload data...`

- 7 Use the Send ASCII File feature of the communications software to send the configuration file to the unit. (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`

- 8 Press any key to return to the configuration menus.

- 9 Reset the Ascend unit, by selecting System > Sys Diag > Sys Reset and confirming the reset.

Restoring passwords

For security reasons, passwords are not written to configuration files created through the serial console. A configuration file created using the Tsave command, however, *does* contain the system passwords. You can restore the Tsave configuration file using the serial console.

After upgrading you may have to re-enter all the passwords on your system. If you edit your saved configuration file, however, and enter passwords in the appropriate fields (by replacing the word **SECURE** in each instance), these passwords will be restored. But note that if you do choose to edit your configuration file, you must save it as text only or you will not be able to load it into your unit.

If you restored a complete configuration, the passwords used in your Security profiles have been wiped out. To reset them:

- 1 Press Ctrl-D to invoke the DO menu, select Password, and choose the Full Access profile.
- 2 When you are prompted to enter the password, press Enter (the null password).
After you have restored your privileges by entering the null password, you should immediately open the Connection profiles, Security profiles, and Ethernet profile (Mod Config menu), and reset the passwords to their previous values.

System messages

Table B-3 explains the messages that can appear during your upgrade.

Table B-3. System software messages

Message	Explanation
UART initialized fat load: bad CRC!! forcing serial download at 57600 bps please download a "thin" system...	The fat load has a CRC (cyclic redundancy check) error. Immediately after this message appears, the serial console speed is switched to 57600 bps, and the Ascend unit initiates an Xmodem serial download. Load a thin load that understand the fat load format, as explained in "Using TFTP to upgrade to a fat or thin load" on page B-6.
File tbam.m40 incompatible fat load format-- discarding downloaded data	You attempted to upgrade to a fat load from a version of system software that does not understand the fat load format. You must first load a thin load that is fat load aware, as explained in "Using TFTP to upgrade to a fat or thin load" on page B-6.
This load has no platform identifier. Proceed with caution.	This message can occur if you are running software version 5.0Ai11 or later and you load an earlier incremental or patch release onto your system. The message indicates that Tloadcode cannot determine which platform the code is intended for. If you are using the correct software version, you can ignore this message.
This load appears not to support your network interface. Download aborted. Use 'tloadcode -f' to force.	Indicates you are attempting to load a version of code intended for a different network interface (for example, loading MAX 4000 T1 software onto a MAX 4000 E1 unit).
This load appears to be for another platform. Download aborted. Use 'tloadcode -f' to force.	Indicates you are attempting to load a version of code onto a platform for which it is not intended (for example, loading MAX 4000 software onto a MAX 2000). This is not recommended
UART initialized fat load: inflate starting system...	Indicates you have successfully loaded a fat load.

Message	Explanation
<pre> UART initialized hybrid load: inflate essential .++. . invalid CRC!! entering restricted mode starting system... </pre>	Indicates the extended load has failed and that your system is being brought up in restricted mode. You must reload the software as explained in “Upgrading software with an extended load” on page B-9.
<pre> UART initialized hybrid load: inflate essential .++. . invalid length!! entering restricted mode starting system... </pre>	Indicates the extended load has failed and that your system is being brought up in restricted mode. You must reload the software as explained in “Upgrading software with an extended load” on page B-9.
<pre> UART initialized hybrid load: inflate essential .++. . inflate expendable starting system... </pre>	Indicates you have successfully loaded an extended load.
<pre> UART initialized thin load: inflate starting system... </pre>	Indicates you have successfully loaded a thin load.

Warranties and FCC regulations

Product warranty C-1

FCC Part 15 C-2

FCC Part 68 Notice C-2

IC CS-03 Notice..... C-3

Product warranty

- 1 Ascend Communications, Inc. warrants that the MAX will be free from defects in material and workmanship for a period of twelve (12) months from date of shipment.
- 2 Ascend Communications, Inc. shall incur no liability under this warranty if
 - the allegedly defective goods are not returned prepaid to Ascend Communications, Inc. within thirty (30) days of the discovery of the alleged defect and in accordance with Ascend Communications, Inc.'s repair procedures; or
 - Ascend Communications, Inc.'s tests disclose that the alleged defect is not due to defects in material or workmanship.
- 3 Ascend Communications, Inc.'s liability shall be limited to either repair or replacement of the defective goods, at Ascend Communications, Inc.'s option.
- 4 Ascend Communications, Inc. MAKES NO EXPRESS OR IMPLIED WARRANTIES REGARDING THE QUALITY, MERCHANTABILITY, OR FITNESS FOR A PARTICULAR PURPOSE BEYOND THOSE THAT APPEAR IN THE APPLICABLE Ascend Communications, Inc. USER'S DOCUMENTATION. Ascend Communications, Inc. SHALL NOT BE RESPONSIBLE FOR CONSEQUENTIAL, INCIDENTAL, OR PUNITIVE DAMAGE, INCLUDING, BUT NOT LIMITED TO, LOSS OF PROFITS OR DAMAGES TO BUSINESS OR BUSINESS RELATIONS. THIS WARRANTY IS IN LIEU OF ALL OTHER WARRANTIES.

Warranty repair

- 1 During the first three (3) months of ownership, Ascend Communications, Inc. will repair or replace a defective product covered under warranty within twenty-four (24) hours of receipt of the product. During the fourth (4th) through twelfth (12th) months of ownership, Ascend Communications, Inc. will repair or replace a defective product covered under warranty within ten (10) days of receipt of the product. The warranty period for the replaced product shall be ninety (90) days or the remainder of the warranty period of the original unit, whichever is greater. Ascend Communications, Inc. will ship surface freight. Expedited freight is at customer's expense.

- 2 The customer must return the defective product to Ascend Communications, Inc. within fourteen (14) days after the request for replacement. If the defective product is not returned within this time period, Ascend Communications, Inc. will bill the customer for the product at list price.

Out-of warranty repair

Ascend Communications, Inc. will either repair or, at its option, replace a defective product not covered under warranty within ten (10) working days of its receipt. Repair charges are available from the Repair Facility upon request. The warranty on a serviced product is thirty (30) days measured from date of service. Out-of-warranty repair charges are based upon the prices in effect at the time of return.

FCC Part 15



Warning: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to Part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio frequency energy, and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference in which case the user will be required to correct the interference at his or her own expense.

The authority to operate this equipment is conditioned by the requirement that no modifications will be made to the equipment unless the changes or modifications are expressly approved by Ascend Communications, Inc.

FCC Part 68 Notice

This Ascend equipment complies with Part 68 of the FCC rules. Located on the equipment is a label that contains, among other information, the FCC registration number. If requested, this information must be provided to the telephone company.

This equipment cannot be used on the telephone company-provided coin service. Connection to Party Line Service is subject to State Tariffs.

If this equipment causes harm to the telephone network, the telephone company will notify you in advance that temporary discontinuance of service may be required. If advance notice isn't practical, the telephone company will notify the customer as soon as possible. Also, you will be advised of your right to file a complaint with the FCC if you believe it is necessary.

The telephone company may make changes in its facilities, equipment, operations, or procedures that could affect the operation of the equipment. If this happens, the telephone company will provide advance notice in order for you to make the necessary modifications in order to maintain uninterrupted service.

If trouble is experienced with this equipment, please contact:

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

If the trouble is causing harm to the telephone network, the telephone company may request you to remove the equipment from the network until the problem is resolved.

It is recommended that the customer install an AC surge arrestor in the AC outlet to which this device is connected. This is to avoid damage to the equipment caused by local lightening strikes and other electrical surges.

This equipment uses the following USOC jacks and codes:

Model Name	Facility Interface Code	Service Order Code	Jack Type
04DU9-BN	P220-T1	6.0N	RJ48C
04DU9-DN	P220-T1	6.0N	RJ48C
04DU9-1KN	P220-T1	6.0N	RJ48C
04DU9-1SN	P220-T1	6.0N	RJ48C
04DU9-1ZN	P220-T1	6.0N	RJ48C

IC CS-03 Notice

The Industry Canada label identifies certified equipment. This certification means that the equipment meets certain telecommunications network protective, operational and safety requirements as prescribed in the appropriate Terminal Equipment Technical Requirements document(s). The Department does not guarantee the equipment will operate to the user's satisfaction.

Before installing this equipment, users should ensure that it is permissible to be connected to the facilities of the local telecommunications company. The equipment must also be installed using an acceptable method of connection. The customer should be aware that compliance with the above conditions may not prevent degradation of service in some situations.

Repairs to certified equipment should be coordinated by a representative designated by the supplier. Any repairs or alterations made by the user to this equipment, or equipment malfunctions, may give the telecommunications company cause to request the user to disconnect the equipment.

Users should ensure for their own protection that the electrical ground connections of the power utility, telephone lines and internal metallic water pipe system, if present, are connected together. This precaution may be particularly important to rural areas.



Caution: Users should not attempt to make such connections themselves, but should contact the appropriate electric inspection authority, or electrician, as appropriate.

Glossary

10Base-T—The 802.3 IEEE standard for operating a 10-Mbps Ethernet network with twisted-pair cabling and a wiring hub. 10BaseT is also known as *UTP Ethernet* and *twisted-pair Ethernet*. See also *10Base-T hub*.

10Base-T hub—A hub providing a common termination point for hosts connected to 10BaseT wiring. See also *10Base-T*.

10-Mbps Ethernet card—A Pipeline 220 contains one 10Base-T Ethernet interface. It provides full 10-Mbps access to an Ethernet network. See also *10Base-T*.

802.2—An IEEE protocol specification for the Media Access Control (MAC) header of an IPX frame in NetWare 3.12 or later. An 802.2 frame contains the Logical Link Control (LLC) header in addition to the MAC header. Compare with *802.3*, *Ethernet II*, *SNAP*. See also *IPX frame*, *LLC*, *MAC*.

802.3—An IEEE protocol specification for the Media Access Control (MAC) header of an IPX frame in NetWare 3.11 or earlier. An 802.3 frame does not contain the Logical Link Control (LLC) header in addition to the MAC header. The 802.3 frame is also called *Raw 802.3*. Compare with *802.2*, *Ethernet II*, *SNAP*. See also *IPX frame*, *LLC*, *MAC*.

802.5—An IEEE protocol specification for the physical layer and Media Access Control (MAC) sublayer of a token-ring topology. 802.5 implements token passing over Shielded Twisted Pair (STP) cabling, and offers data rates of 4 or 16 Mbps. See also *STP cable*.

ABR—Area Border Router. An ABR is an Open Shortest Path First (OSPF) router that belongs to both a regular area and the backbone area. See also *area*, *backbone area*, *OSPF*.

address resolution—A method of mapping a logical address (such as an IP address) to a hardware address (such as a MAC address). See also *ARP*, *hardware address*, *IP address*, *logical address*, *MAC address*.

Address Resolution Protocol—See *ARP*.

adjacency—A relationship formed between neighboring Open Shortest Path First (OSPF) routers for the purpose of exchanging routing information. An OSPF router dynamically detects its neighboring routers by sending Hello packets to the multicast address AllSPFRouters. It then attempts to form adjacencies. Neighbors exchange databases and build a consistent, synchronized database between them.

When an OSPF router detects a change on one of its interfaces, it modifies its link-state database and multicasts the change to its adjacent neighbor, which in turn propagates the change to its adjacent neighbor, until all routers within an area have synchronized link-state databases. This method of updating routing information results in quick convergence among routers.

See also *area*, *convergence*, *link-state database*, *OSPF*, *router*.

agent—A network device (such as the Pipeline 220) that provides Simple Network Management Protocol (SNMP) information to a manager application running on another computer. The agent and manager share a database of information, called the *Management Information Base (MIB)*. The manager polls the agent for information on at regular intervals. When an unusual system event occurs, the agent can use a message called a *traps-PDU* to send unsolicited information to the manager. See also *manager*, *MIB*, *SNMP*, *traps-PDU*.

American National Standards Institute—See *ANSI*.

ANSI—American National Standards Institute. ANSI creates standards for networking and communications. It is the U.S. representative to the International Standards Organization (ISO). See also *ISO*.

Answer-Defaults profile—A profile that sets baseline values to determine how the Pipeline 220 evaluates incoming calls before it accepts them. If the call does not comply with the Answer-Defaults settings, the unit rejects the call without answering it. Therefore, you must check the Answer-Defaults values to make sure they are appropriate for your site.

The Pipeline 220 applies the Answer-Defaults values before it routes the call to a digital modem or High-Level Data Link Control (HDLC) channel for processing, and before it locates a Connection profile. If the caller's profile contains a parameter or attribute similar to one in the Answer-Defaults profile, but the caller's setting specifies a different value, the Pipeline 220 uses the value in the Connection profile to build the session.

By default, the Answer-Defaults profile enables all types of encapsulation and routing, and the basic call-setup parameters use the lowest-common-denominator settings. The default settings are appropriate for many sites. You might want to change the settings in order to finetune the criteria by which the Pipeline 220 accepts calls or determines how much bandwidth is accessible to Multilink Protocol (MP) or Multilink Protocol Plus (MP+) sessions.

See also *Connection profile*, *digital signal*, *HDLC channel*.

AppleTalk—Apple's protocol suite that enables Macintosh computers to function on a network. AppleTalk works with such network operating systems as TOPS (from Sun Microsystems), NetWare for Macintosh (from Novell), and AppleShare.

Application layer—The highest layer of the OSI Reference Model. The Application layer provides applications with access to the network. File transfer, email, and network management software are examples of Application-layer programs. Protocols such as File Transfer Protocol (FTP), Rlogin, Simple Network Management Protocol (SNMP), and Telnet provide Application-layer services. See also *FTP*, *OSI Reference Model*, *SNMP*, *Telnet*.

area—A portion of an Open Shortest Path First (OSPF) Autonomous System (AS). An area acts as its own network. All area-specific routing information stays within the area, all routers within an area have a synchronized link-state database, and each database within an area is unique. On the Pipeline 220, an area number uses dotted decimal notation. It is not an IP address. To tie the areas together, some routers belong to a backbone area and one other type of area. These routers are called *Area Border Routers (ABRs)*. See also *ABR*, *AS*, *backbone area*, *link-state database*, *normal area*, *NSSA*, *OSPF*, *router*, *stub area*.

Area Border Router—See *ABR*.

ARP—Address Resolution Protocol. ARP is a protocol in the TCP/IP protocol suite. By mapping an IP address to a physical (hardware) address, ARP enables a unit to identify hosts on an Ethernet LAN. See also *Ethernet, proxy ARP, TCP/IP*.

AS—Autonomous System. An AS is a group of Open Shortest Path First (OSPF) routers that exchange information, typically under the control of one company. An AS can include a large number of networks, all of which share the same AS number. All information exchanged within the AS is interior. Exterior protocols, such as Exterior Gateway Protocol (EGP), exchange routing information between one AS and another. Using an EGP, the Pipeline 220 imports external routes into its OSPF database and flags them as ASE (Autonomous System External). See also *ASE, EGP, external route, OSPF, router*.

Ascend Tunnel Management Protocol—See *ATMP*.

ASCII—American Standard Code for Information Interchange. ASCII is a character-encoding system for Local Area Networks (LANs). The 128 standard ASCII characters are composed of seven bits, and have the values 0–127. The extended ASCII character set contains another 128 values.

ASE—Autonomous System External. The Pipeline 220 uses the term ASE to denote external routes it imports into its Open Shortest Path First (OSPF) database. The Pipeline 220 redistributes these routes by means of OSPF ASE advertisements, and propagates its OSPF routes to remote WAN routers running Routing Information Protocol (RIP). See also *external route, OSPF, RIP, router*.

ASE Type-5—Autonomous System External Type-5. ASE Type-5 is an external route originated by an Area Border Router (ABR) as a Link-State Advertisement (LSA). An Open Shortest Path First (OSPF) normal area allows Type-5 LSAs to be flooded throughout it. A Not So Stubby Area (NSSA) and a stub area do not receive or originate Type-5 LSAs. Compare with *ASE Type-7*. See also *ABR, AS, ASE, external route, LSA, normal area, NSSA, OSPF, stub area*.

ASE Type-7—Autonomous System External Type-7. ASE Type-7 is a type of Link-State Advertisement (LSA) defined for Not So Stubby Areas (NSSAs) in Open Shortest Path First (OSPF) version 2. The Pipeline 220 can import ASE Type-7s only from static route definitions. Only a single NSSA can originate and advertise a Type-7 LSA. Type-7 LSAs are not flooded throughout the area like Type-5 LSAs. Compare with *ASE Type-5*. See also *AS, ASE, LSA, NSSA, OSPF, stub area*.

ASN.1—Abstract Syntax Notation One. In the OSI Reference Model, ASN.1 is a notation for describing data structures on a network. It provides a consistent syntax for transferring data between different systems. See also *OSI Reference Model*.

ATM—Asynchronous Transfer Mode. ATM is a packet-switched, broadband network architecture central to Broadband ISDN (B-ISDN). It provides very high bandwidth, enabling data, voice, and multimedia transmissions to occupy the same line. ATM is also known as *cell relay*. See also *, packet switching*.

ATMP—Ascend Tunnel Management Protocol. ATMP provides a tunneling mechanism between two Ascend units across the Internet or a Frame Relay network. Each Ascend unit can be a Pipeline 220 or a Pipeline 400. The protocol uses standard Generic Routing Encapsulation (GRE) and is based on the User Datagram Protocol (UDP) and Internet Protocol (IP).

ATMP provides a Virtual Private Network (VPN) solution over the backbone resources of Internet Service Providers (ISPs) and carriers. Without ATMP, each mobile node and remote user has to dial directly into the network, resulting in long-distance charges. With ATMP, users can make a local call and have the transmission securely tunneled.

As described in RFC 1701, GRE hides packet contents and enables transmission of packets that the Internet would otherwise not accept. When you use ATMP with the Pipeline 220, you can transmit either IP packets that use unregistered addresses or IPX packets from roaming clients.

See also *Frame Relay, GRE, IP, IPX, ISP, UDP, VPN*.

AUI—Auxiliary Unit Interface. An AUI is a 15-pin D-type connector for Ethernet connections. It typically links a cable to a Network Interface Card (NIC). An AUI is also known as a *Digital, Intel, Xerox (DIX) connector*. See also *Ethernet, NIC*.

authentication—A method of identifying users permitted to access network resources. Authentication is the first line of defense against unauthorized access to your network. The Pipeline 220 supports the following authentication method:

- Name and password authentication of Point-To-Point Protocol (PPP) calls. The Pipeline 220 supports Password Authentication Protocol (PAP), PAP with encryption (PAP-DES), Challenge Handshake Authentication Protocol (CHAP), and Microsoft's extension of CHAP (MS-CHAP).

When the Pipeline 220 is shipped from the factory, it is set to not require any authentication. See also *CHAP, PAP*.

authorization—Permission for a user to carry out a certain set of tasks after he or she has access to your LAN. Authorization occurs *after* authentication is complete. On the Pipeline 220, you configure authorization in the following profiles:

- The Terminal-Server profile, to restrict access to the terminal-server software.
- The SNMP profile, to restrict access to the system by means of Simple Network Management Protocol (SNMP) manager utilities.
- System-wide and Connection profiles, to restrict access to certain Domain Name System (DNS) servers.

See also *authentication*.

Autonomous System—See *AS*.

Auxiliary Unit Interface—See *AUI*.

backbone—The part of the communications network designed to carry the bulk of the traffic. The backbone provides connectivity between subnets in an enterprise-wide network. See also *enterprise-wide network, IP subnet*.

backbone area—An Open Shortest Path First (OSPF) area that connects routers for the purpose of hierarchical routing. The backbone area is special and always has the area number 0.0.0.0. To tie areas together, some routers belong to the backbone area and one other area. These routers are called *Area Border Routers (ABRs)*. See also *ABR, area, OSPF, router*.

backbone network—A network with a central cabling scheme linking it to other networks. Hosts on networks linked to the backbone can communicate with one another.

backbone router—A router attached to a backbone network by nailed-up lines. Usually, a backbone router does not have any built-in digital dial-up WAN interfaces. Manufacturers of backbone routers include Cisco, Wellfleet, 3Com, and CrossCom. See also *backbone network, router*.

Backup Designated Router—See *BDR*.

Backward Explicit Congestion Notification—See *BECN*.

bandwidth—The amount of data a link can carry, measured bits per second (bps) for digital signals, and in hertz (Hz) for analog signals. See also *digital signal*.

BDR—Backup Designated Router. A BDR is the router that the Open Shortest Path First (OSPF) area uses in the event that the Designated Router (DR) goes out of service. To prevent the DR from becoming a serious liability to the network if it fails, OSPF elects a Backup Designated Router (BDR). Other routers maintain adjacencies with both the DR and BDR, but the backup router leaves as many processing tasks as possible to the DR. If the DR fails, the backup immediately becomes the DR and a new backup is elected.

The Pipeline 220 can function as either a DR or a BDR. However, many sites choose to assign LAN-based routers to these functions in order to dedicate the Pipeline 220 to WAN processing. See also *adjacency, area, DR, OSPF, router*.

bearer service—An ISDN service for transmitting information from one device to another. Common bearer services are circuit-switched, Frame Relay, and X.25 services. See also *Frame Relay*.

BECN—Backward Explicit Congestion Notification. BECN is a bit set in a Frame Relay header to notify a source node that there is traffic congestion on the network. See also *FECN, Frame Relay*.

BGP—Border Gateway Protocol version 4. BGP routes packets between networks that use different types of protocols. It is known as an *Exterior Gateway Protocol (EGP)*, and replaces an older protocol called EGP. See also *EGP*.

Bit—Binary digit, the smallest unit of information a computer can process, representing one of two states (indicated by 1 and 0).

bits per second—See *bps*.

black-hole interface—An interface that enables the router to handle packets whose IP address matches an unused IP address in a summarized address pool. The black-hole interface has an IP address of 127.0.0.3. When you specify this address as the router to the destination pool network, the Pipeline 220 silently discards packets to an invalid host on that network. See also *POP*.

BOOTP—Boot Protocol. BOOTP starts up a network device by using information from a server. The Pipeline 220 can use BOOTP to get settings and check for a new software load. In addition, you can enable the terminal server to respond to BOOTP within a Serial Line Internet Protocol (SLIP) session. An interactive user who initiates a SLIP session can get an IP address from a designated IP address pool by means of BOOTP. See also *IP address, terminal server*.

Boot Protocol—See *BOOTP*.

Border Gateway Protocol version 4—See *BGP*.

bps—A nested acronym, meaning binary digits per second, and a measure of the capacity of a line.

bridge—A hardware device that transmits packets between networks. A bridge forwards packets from one network to another, and discards packets destined for hosts on the sending network. Operating at the Data Link layer, a bridge makes multiple networks look like a single network to higher-level protocols and software. See also *Data Link layer*.

bridge entry—An entry in a bridging table. The Pipeline 220 is a transparent bridge (also called a *learning bridge*). As the Pipeline 220 forwards a packet, it notes the packet's source address and creates a bridging entry that associates a host's Media Access Control (MAC) address with a particular Ethernet interface. The Pipeline 220 also learns about bridging links from Connection profiles, either because the remote caller used the profile to dial the link, or because the profile matched an incoming call. In addition, you can specify static bridge entries in a local profile. See also *bridge*, *bridging table*, *Connection profile*, *Ethernet*, *MAC*.

bridging—A method of moving packets between networks by means of a device called a *bridge*, which operates at the Data Link layer. See also *bridge*, *Data Link layer*.

bridging table—A table that contains entries pairing up a host's Media Access Control (MAC) address with a particular Ethernet interface. If the Pipeline 220 receives a packet whose destination MAC address is not on the local network, it first checks its bridging table. If it find the packet's destination MAC address, the Pipeline 220 dials the connection and bridges the packet. If it does not find the address, the Pipeline 220 checks for active sessions that have bridging enabled. If one or more active bridging links are up, the Pipeline 220 forwards the packet across all active sessions that have bridging enabled. See also *bridge*, *bridge entry*, *Ethernet*, *MAC*.

broadcast network—A network in which the router sends packets to all users, whether they appear on subscription lists or not. In an Open Shortest Path First (OSPF) topology, a broadcast network is any network that has more than two OSPF routers attached and can address a single physical message to all of them. See also *OSPF*, *multicast network*, *router*, *unicast network*.

bus—A path for signals transmitted between a computer's CPU and other hardware devices.

byte—8 bits of data, also called an *octet*.

call—A single session in which a calling device and an answering device connect over the WAN.

Call Detail Reporting—See *CDR*.

CCITT—Consultative Committee on International Telegraphy and Telephony. The CCITT is a disbanded organization whose standards were moved to the UN-sanctioned ITU-T on March 1, 1993.

CCP—Compression Control Protocol. CCP enables both ends of a Point-To-Point Protocol (PPP) connection to negotiate whether to use data compression, and if so, which algorithm to use.

CDR—Call Detail Reporting. CDR is a feature that provides a database of information about each call, including date, time, duration, called number, calling number, call direction, service type, and associated inverse multiplexing session and port. Because the network carrier charges for bandwidth on an as-used basis, and bills each connection in an inverse-multiplexed call as a separate charge, you can use CDR to understand and manage bandwidth usage and the cost of each inverse-multiplexed session.

You can manipulate CDR information in order to create a wide range of different reports, including reports based on individual call costs, inverse-multiplexed WAN session costs, costs on an application-by-application basis, and bandwidth-usage patterns over specified time periods. You can use this information to better understand your bandwidth usage patterns and, if necessary, make adjustments to the ratio of switched to dedicated bandwidth between network sites.

Central Office—See *CO*.

Central Processing Unit—See *CPU*.

Challenge Handshake Authentication Protocol—See *CHAP*.

channel—A portion of a line's bandwidth. A line contains a fixed number of channels. Each line can contain switched channels only, nailed-up channels only, or a combination of switched and nailed-up channels. See also *bandwidth*, *line*, *nailed-up channel*.

channelized T1 PRI/E1 PRI—A T1 PRI or E1 PRI line divided into individual 64-Kbps channels, or into channels whose data rate is a multiple of 64 Kbps (such as a 256-Kbps channel made from four 64-Kbps channels). Channelized T1 PRI or E1 PRI lines can consist of switched lines with inband signaling, or nailed-up lines. For example, a nailed-up line can run from the Central Office (CO) to the corporate headquarters as a single, unchannelized T1 PRI or E1 PRI line, and can then be divided into channels when it runs to remote sites from the corporate headquarters. See also *E1 line*, *E1 PRI line*, *nailed-up line*, *T1 line*, *T1 PRI line*, *unchannelized service*.

Channel Service Unit—See *CSU*.

CHAP—Challenge Handshake Authentication Protocol. CHAP authentication verifies the caller's identity by using a three-way handshake upon initial link establishment, and then by repeating the handshake any number of times. In CHAP authentication, the authentication server sends a challenge to the caller. The caller responds with an MD5 digest calculated from the password. The authentication server then checks the digest against its own calculation of the expected hash value to authenticate the call. The server can send a new challenge at random intervals.

CHAP is a stronger authentication method than Password Authentication Protocol (PAP), because the password does not travel across the line as plain text. In addition, the use of repeated challenges limits the time of exposure to any single attempt to break the encryption code, and the server is in control of how often it sends challenges. See also *encryption*, *PAP*.

circuit—A connection between endpoints over a physical medium.

circuit connection—A connection that follows a specified path through the Frame Relay switch. By linking two Data Link Connection Indicator (DLCI) endpoints, the Pipeline 220 creates a Permanent Virtual Circuit (PVC). The two DLCI endpoints act as a tunnel. Data that the Pipeline 220 receives on one DLCI bypasses the Ascend router and goes out on the other

DLCI. If any one of the DLCIs in a PVC becomes inactive because of a disconnect or failure, the PVC using that DLCI becomes inactive. A physical line can carry multiple DLCIs, and the failure of the line causes the failure of all the DLCIs it carries. Compare with *gateway connection*, *redirect connection*. See also *DLCI*, *Frame Relay switch*, *PVC*, *router*.

client DNS—A configuration that enables the Pipeline 220 to direct incoming connections to a Domain Name System (DNS) server belonging to a particular client or location, thereby preventing WAN users from accessing a local DNS server. Client DNS has two levels: a global configuration that applies to all PPP connections, and a connection-specific configuration. The Pipeline 220 uses the global client addresses only if none are specified in the Connection profile. The addresses configured for client DNS servers are presented to WAN connections during IPCP negotiation. You can also choose to present your local DNS servers if no client servers are defined or available. See also *Domain Name System*, *IPCP*.

clock—A timing mechanism for synchronizing data communication and processing tasks. A clock divides time into very short intervals. The clock speed is the number of intervals per second. See also *clock source*.

clock source—The master source for clocking of synchronous connections. The entire multishelf Pipeline 220 system uses a single synchronous clock source. The Pipeline 220 chooses the clock source from the T1 or E1 lines you specify as possible external sources. If there are no eligible external sources, the system uses an internal clock generated by the master shelf controller.

You can use the Clock-Source diagnostic command to determine the current master clock source. If you execute the command on the shelf controller, the output tells which slot (if any) is the clock source. If you execute the command on a T1 or E1 card, the output tells which line is the clock source.

See also *clock*, *synchronous transmission*.

CO—Central Office. The CO is the telephone switching office to which a customer directly connects. It connects the customer to other portions of the telephone network.

command mode—A terminal-server mode in which you can enter commands at the terminal-server prompt. Compare with *menu mode*.

community name—A password that the Pipeline 220 sends to the Simple Network Management Protocol (SNMP) manager when an SNMP trap event occurs, and that the manager sends to the Pipeline 220 with each polling request. The password authenticates the sender. The default is `public`. See also *agent*, *manager*, *SNMP*.

compression—A process that removes waste and redundancy in a data file, enabling faster throughput. The results of compression depend on the content of the file. Some files contain a lot of waste, and others contain almost none. See also *VJ compression*.

Connection profile—A local profile containing authentication and configuration information about a remote device or user.

Consultative Committee on International Telegraphy and Telephony

—See *CCITT*.

context—A subprofile that resides directly below a profile, or that is nested within another subprofile. See also *profile*.

convergence—The time it takes all routers to receive information about a change to the network topology. A slow convergence can result in routing loops and errors. A Routing Information Protocol (RIP) router broadcasts its entire routing table every 30 seconds. On a 15-hop network, convergence can be as high as 7.5 minutes. In contrast, Open Shortest Path First (OSPF) uses a link-state database of the network, and propagates only changes to the database, resulting in faster convergence. See also *link-state database*, *OSPF*, *RIP*.

cost—An Open Shortest Path First (OSPF) value you assign to the output side of each router interface. The cost indicates the likelihood that the Pipeline 220 will use the interface to transmit data. The lower the cost, the more likely that the Pipeline 220 will use the interface. You can use the cost to perform preferred path selection. If two paths to a destination have equal costs, you can assign a higher cost to one of the paths, making it a backup when the primary path is not available. In addition, you may want to reflect the bandwidth of a connection when assigning costs. For example, the cost of a single B-channel connection could be 24 times greater than the cost of a T1 link.

The Pipeline 220 has a default cost of 1 for a connected route (Ethernet) and 10 for a WAN link. Be careful when assigning costs. Incorrect cost metrics can cause delays and congestion on the network.

See also *OSPF*, *route*, *router*.

CPE—Customer Premises Equipment. CPE is equipment connected to the telephone network, and located at the customer's site. The equipment can be owned or leased.

CPU—Central Processing Unit. The CPU is the computer's main processor.

CRC—Cyclic Redundancy Check. CRC is an error-detection method that uses a mathematical divisor to check the integrity of the data in a transmitted packet.

crossover cable—A cable with wires that cross over, so that the terminating ends of the cable have opposite wire assignments. Compare with *straight-through cable*.

CSU—Channel Service Unit. Along with a Data Service Unit (DSU), a CSU is a component of Data Circuit-Terminating Equipment (DCE). A CSU connects a digital phone line to a customer's network-access equipment. It can be built into the network interface of the network-access equipment, or it can be a separate device. The CSU terminates the connection at the user's end and processes digital signals. It also prevents a faulty DSU from interfering with data transmissions on the digital line. See also *DCE*, *digital signal*, *DSU*.

Customer Premises Equipment—See *CPE*.

Cyclic Redundancy Check—See *CRC*.

D4-framed T1 line—A T1 line that uses the D4 format, also known as the *Superframe format*, to frame data at the physical layer. The D4 format consists of 12 consecutive frames, each one separated by framing bits. T1 lines that do not use ISDN D-channel signaling use the D4 format. See also *T1 line*.

Database-Description packet—A Type-2 Open Shortest Path First (OSPF) packet. OSPF routers exchange Database-Description packets when an adjacency is being initialized. Each

packet describes the contents of the link-state database. The routers use a poll-response procedure. One of the routers is the master, and the other a slave. The master sends Database-Description poll packets, and the slave sends Database-Description response packets. OSPF links the responses to the polls by means of a sequence number in each packet. See also *adjacency*, *link-state database*, *OSPF*.

Data Circuit-Terminating Equipment—See *DCE*.

data compression—See *compression*.

Data Encryption Standard—See *DES*.

data filter—A packet filter that defines which packets the Pipeline 220 can transmit on a connection. When you apply a data filter, its forward or drop action affects the actual data stream by preventing certain packets from reaching the Ethernet from the WAN, or vice versa. Many sites use data filters for security purposes, but you can apply data filters to any purpose that requires the Pipeline 220 to drop or forward only 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 users to access only specific devices across the WAN. See also *packet filter*.

Datagram Delivery Protocol—See *DDP*.

Data Link Connection Indicator—See *DLCI*.

Data Link layer—The second layer of the OSI Reference Model. The Data Link layer creates, sends, and receives data packets appropriate for the type of network in use. Data-Link-layer protocols include High-Level Data Link Control (HDLC), Link Access Procedure, Balanced (LAPB), Link Access Procedure, D channel (LAPD), Point-To-Point Protocol (PPP), and Serial Line Internet Protocol (SLIP). See also *HDLC*, *LAPD*, *OSI Reference Model*, *PPP*.

Data Service Unit—See *DSU*.

Data Terminal Equipment—See *DTE*.

DCE—Data Circuit-Terminating Equipment. A DCE is a device that connects Data Terminal Equipment (DTE) to a communications channel, such as a telephone line. A DTE refers to a device that an operator uses, such as a computer or a terminal. A DCE converts the format of the data coming from the DTE into a signal suitable to the communications channel. An example of a DCE is a modem, which converts digital data from a computer to analog signals suitable for sending over a telephone line. See also *digital data*, *DTE*.

DDP—Datagram Delivery Protocol. DDP is an AppleTalk Network-layer protocol. It provides connectionless service between sockets, and handles both addressing and routing. See also *routing*, *socket*.

DE—Discard Eligibility. DE is a bit in a Frame Relay packet header. You set the DE bit to indicate that the network can discard the packet when traffic reaches a high level. See also *Frame Relay*.

default gateway—The default router the Ascend unit uses for traffic from a specific connection if it finds no explicit route in the IP routing table. See also *IP router*, *IP routing table*.

default route—The route the Ascend unit uses if it does not find a match for a packet's destination address. The default route has the destination address 0.0.0.0. If the Ascend unit finds a default route, it brings up the required connection (if necessary) and forwards the packet. If the routing table has no default route and no route that matches a packet's destination address, the Pipeline 220 drops the packet. See also *IP route*, *IP router*, *IP routing table*.

DES—Data Encryption Standard. DES is the U.S. encryption standard for nonclassified documents. This standard uses a 64-bit key and private-key encryption. In private-key encryption, only the sender and receiver know the key for encrypting the data. DES cannot ensure that the sender and receiver are legitimate. A sender who has learned the key can fraudulently use it. See also *encryption*.

Designated Router—See *DR*.

Destination Service Access Point—See *DSAP*.

DHCP—Dynamic Host Configuration Protocol. DHCP is a TCP/IP protocol that enables a client to obtain a temporary IP address from a central server (known as a *DHCP server*). See also *DHCP server*.

DHCP server—Dynamic Host Configuration Protocol server. A DHCP server assigns a temporary IP address to a client that requests it. See also *DHCP*, *DHCP spoofing*, *IP address*.

DHCP spoofing—Dynamic Host Configuration Protocol spoofing. A process that enables a local device to receive an IP address from a DHCP server across a slow WAN link. Typically, a device requesting an IP address from a DHCP server waits a limited amount of time before timing out the request. For complex WAN links, with authentication processes, there may not be enough time to complete the process. DHCP spoofing operates as follows:

- 1 The PC sends a broadcast DHCP request.
 - 2 Acting as a DHCP server, the Pipeline receives the DHCP request, and sends the PC a temporary IP address. The address can be static or dynamic. It has a very short time-to-live (ttl).
 - 3 The Pipeline dials the remote side, passing along the original DHCP request.
 - 4 The DHCP server sends back a server-assigned IP address.
 - 5 When the Pipeline receives the address from the remote side, it passes the address to the PC.
 - 6 The PC changes its IP address to the server-assigned address.
- See also *DHCP*, *IP address*.

digital data—Data that can have only a limited number of separate values. The time of day represented by a digital clock, or the temperature represented by a digital thermometer are examples of digital data. The digital values do not change continuously, but remain at one discrete value and then change to another discrete value. See also *digital signal*.

digital line—A line that transmits data by means of a digital signal. See also *digital signal*.

digital signal—A type of signal that uses a limited number of discrete values to encode data transmitted over a wire. The value of the data encoded in a digital signal depends upon the state of the signal during a particular time period. Therefore, the sender and the receiver must synchronize their clocks. Each clock runs at a baud rate, the number of times per second the state of the signal is read or set. Several clocking schemes are available, and digital signals often include clock timing cues. A digital signal can transmit analog or digital data. For example, a Compact Disc (CD) encodes music data into digital signals, while the wires between computers transmit digital data in digital signals. See also *digital data*.

Digital Signal Cross-Connect—See *DSX*.

direct route—A route that can reach a destination without going through any intervening routers. See also *route*, *router*.

Discard Eligibility—See *DE*.

distance-vector metric—A metric that uses a hop count to select the shortest route to a destination network. Routing Information Protocol (RIP) always uses the lowest hop count, regardless of the speed or reliability of a link. Compare with *link-state metric*. See also *RIP*.

DIX connector—See *AUI*.

DLCI—Data Link Connection Indicator. A DLCI is a number between 16 and 991 that the Frame Relay administrator assigns. It identifies a logical link between a device and a Frame Relay switch. A Connection profile specifies a DLCI for each user connection. The Frame Relay switch uses the DLCI to route frames through the network. The DLCI can change as frames pass through multiple switches. See also *Connection profile*, *Frame Relay switch*.

DNS—Domain Name System. DNS is a TCP/IP service for centralized management of address resolution. Using DNS, you can specify a symbolic name instead of an IP address. A symbolic name consists of a user name and a domain name in the format *username@domain_name*. The user name corresponds to the host number in the IP address. The domain name corresponds to the network number in the IP address. A symbolic name might be *steve@abc.com* or *joanne@xyz.edu*. The domain identifier is the last part of the domain name, and identifies the type of organization to which the host belongs. DNS maintains a database of network numbers and corresponding domain names. When you use a symbolic name, DNS translates the domain name into an IP address, and sends it over the network. When the Internet service provider receives the message, it uses its own database to look up the user name corresponding to the host number. See also *address resolution*, *host number*, *IP address*, *network number*.

domain identifier—The portion of a domain name that appears last and specifies the type of organization to which the host belongs. The Internet Network Information Center (InterNIC) specifies the following domain identifiers:

Domain identifier	Description
.arpa	ARPANET
.com	Commercial enterprise
.edu	Educational institution
.gov	Governmental organization

Domain identifier	Description
.mil	Military organization
.org	An organization not covered by the other categories

domain name—The portion of a symbolic name that corresponds to the network number in the IP address. In the symbolic name *steve@abc.com*, the domain name is *abc.com*. See also *IP address*, *network number*.

Domain Name System—See *DNS*.

dotted decimal notation—A system for specifying an IP address or subnet mask. In dotted decimal notation, each of the four portions of the IP address or mask is separated from the others by a decimal point, as in the address *200.10.5.1*. See also *IP address*, *subnet mask*.

DR—Designated Router. The DR is the router with which all other Open Shortest Path First (OSPF) routers in a broadcast network establish adjacencies. To reduce the number of adjacencies each router must form, OSPF calls one of the routers the DR. Doing so simplifies the routing table update procedure and reduces the number of link-state records in the database. The DR plays other important roles in reducing the overhead of OSPF link-state procedures. For example, other routers send Link-State Advertisements (LSAs) to the DR by using the “all-designated-routers” multicast address of *224.0.0.6*.

The administrator chooses the DR based on the processing power, speed, and memory of the system, and then assigns priorities to other routers in case the Backup Designated Router (BDR) is down at the same time. The Pipeline 220 can function as a DR or BDR. However, many sites choose to assign a LAN-based router as the DR or BDR in order to dedicate the Pipeline 220 to WAN processing.

See also *adjacency*, *BDR*, *LSA*, *OSPF*, *router*.

DRAM—Dynamic Random Access Memory. DRAM is a kind of memory whose information resides in capacitors. The charge of each capacitor must be periodically refreshed. Compare with *EEPROM*, *NVRAM*, *RAM*.

DRAM upgrade slot—A slot on the Pipeline 220 shelf controller that enables you to add DRAM upgrades. See also *DRAM*.

DS0 channel—A 64-Kbps D channel on a digital line. See also *DS1*.

DS1—A 1.544-Mbps channel that consists of 24 DS0 channels and an extra framing bit. A DS1 channel uses either the D4 or ESF method of framing. You can transmit DS1 signals over a T1 line. See also *D4-framed T1 line*, *DS0 channel*, *ESF*, *T1 line*.

DS2—A 6.312-Mbps channel that consists of four DS1 channels. See also *DS1*.

DS3—A 44.736-Mbps channel that consists of seven DS2 channels. See also *DS2*.

DSAP—Destination Service Access Point. A DSAP is the Service Access Point (SAP) address at which the Logical Link Control (LLC) layer passes information to a Network-layer process. See also *SAP*, *SSAP*.

DSU—Data Service Unit. Along with a Channel Service Unit (CSU), a DSU is a component of Data Circuit-Terminating Equipment (DCE). The DSU connects to Data Terminal

Equipment (DTE) by means of a synchronous serial interface, such as a V.35, RS-422, or RS-423 connection. The DSU formats and controls the flow of digital data between the network and the CSU. See also *CSU*, *DCE*, *digital data*, *DTE*, *V.35*.

DSX—Digital Signal Cross-Connect. DSX is a method of connecting DS1 and DS3 signals by linking T1 and T3 lines. See also *DS1*, *DS3*, *T1 line*.

DTE—Data Terminal Equipment. A DTE is a device that an operator uses, such as a computer or a terminal. Compare with *DCE*.

Dynamic Host Configuration Protocol—See *DHCP*.

Dynamic Random Access Memory—See *DRAM*.

dynamic route—A path to another network that the router learns by means of dynamic updates from other routers, rather than by means of a static specification in a configured profile. Routers that use Routing Information Protocol (RIP) broadcast their entire routing tables every 30 seconds, updating other routers about which routes are usable. Hosts that run Internet Control Message Protocol (ICMP) can also send ICMP Redirects to offer a better path to a destination network. Open Shortest Path First (OSPF) routers propagate link-state changes as they occur in order to update their routing tables. Compare with *multipath route*, *static route*. See also *IP route*, *IPX route*, *route*.

E1 line—A line that supports 32 64-kbps channels, each of which can transmit and receive data or digitized voice. The line uses framing and signaling to achieve synchronous and reliable transmission. The most common configurations for E1 lines are E1 PRI and unchannelized E1. See also *E1 line*, *E1 PRI line*, *unchannelized service*.

E1 PRI line—E1 Primary Rate Interface line. An E1 PRI line consists of 32 64-Kbps channels. It uses 30 B channels for user data, 1 64-Kbps D channel for ISDN D-channel signaling, and one framing channel. The B channels can be all switched, all nailed up, or a combination of switched and nailed up. The E1 PRI line is a standard in Europe and Asia called CEPT G.703. Compare with *T1 PRI line*, *unchannelized service*. See also *DDP*, *E1 line*, *nailed-up channel*.

E1 Primary Rate Interface line—See *E1 PRI line*.

Echo—A signal that determines whether a node can receive and acknowledge data transmissions. A host sends an Echo packet. If the destination is properly connected and receives the Echo packet, it sends back an Echo Reply.

ECN—Explicit Congestion Notification. ECN is a method of informing Frame Relay nodes that there is traffic congestion on the network. The Frame Relay header can use a Backward Explicit Congestion Notification (BECN) bit or a Forward Explicit Congestion Notification (FECN) bit to notify nodes of traffic congestion. *BECN*, *FECN*, *Frame Relay*.

EEPROM—Electrically Erasable Programmable Read-Only Memory. EEPROM is a type of Programmable Read-Only Memory (PROM) that can be erased by exposing it to an electrical charge. It retains its contents across resets and power cycles, and is similar to NVRAM. With EEPROM, data is written or erased one byte at a time; with NVRAM, data is written or erased in blocks. See also *NVRAM*, *PROM*.

EGP—Exterior Gateway Protocol. EGP is a type of protocol used to exchange routing information between one Open Shortest Path First (OSPF) Autonomous System (AS) and another. The AS number may be used by Area Border Routers (ABRs) to filter out certain EGP routing information. OSPF can make use of EGP data generated by other border routers and added to the OSPF system as Autonomous System Externals (ASEs). See also *ABR*, *AS*, *ASE*, *OSPF*.

EIA—Electronic Industries Association. The EIA is a group that determines standards for electrical transmission.

EIA/TIA-232—A Physical-layer standard nearly identical to V.24. EIA/TIA-232 is also known as *RS-232*. See also *RS-232*.

Electrically Erasable Programmable Read-Only Memory—
See *EEPROM*.

Electronic Industries Association—See *EIA*.

encapsulation—A technique used by layered protocols in which a low-level protocol accepts a message from a higher-level protocol, and then places the message in the data portion of the lower-level frame. The logistics of encapsulation require that packets traveling over a physical network contain a sequence of headers. Encapsulation enables the transmission of data over networks that use differing protocols.

encryption—A process that takes ordinary data and converts it into a format unreadable to anyone without a decryption key. Authorized personnel with access to this key can unscramble the information. Data encryption is a useful tool against network snoopers. See also *public-key encryption*.

enterprise-wide network—A network that contains all or most of a company's hardware and software resources. Typically, an enterprise-wide network includes computers that run different operating systems and reside on different types of networks. Therefore, achieving interoperability is the biggest challenge facing the administrator of an enterprise-wide network.

ESF—Extended SuperFrame. ESF is a framing format that consists of 24 consecutive frames, separated by framing bits. The ISDN specification advises that you use ESF with ISDN D-channel signaling.

Ethernet—The most commonly used architecture for Local Area Networks (LANs), connecting devices such as computers, printers, and terminals. An Ethernet network uses the Physical and Data Link layers for data transmission. Ethernet incorporates a bus topology, and can operate at a rate of up to 10 Mbps. See also *Data Link layer*, *Physical layer*.

Ethernet II—A protocol specification for the Media Access Control (MAC) header of an IPX frame. Compare with *802.2*, *802.3*, *SNAP*. See also *IPX frame*, *MAC*.

Ethernet card—A card that provides the Pipeline 220 with access to Ethernet networks. See also *10-Mbps Ethernet card*.

Ethernet transceiver—A device that connects workstations to standard thick or thin Ethernet-style cable. An Ethernet transceiver sends and receives information, and offers data-packet collision detection. See also *Thick Ethernet*, *Thin Ethernet*.

EU-RAW—A WAN encapsulation protocol used primarily in Europe. IP packets are HDLC-encapsulated and include a Cyclic Redundancy Check (CRC).

EU-UI—A WAN encapsulation protocol used primarily in Europe. IP packets are HDLC-encapsulated, and include a special header and a Cyclic Redundancy Check (CRC).

Explicit Congestion Notification—See *ECN*.

Extended SuperFrame—See *ESF*.

Exterior Gateway Protocol—See *EGP*.

external route—A route imported into the Open Shortest Path First (OSPF) database from outside the router's Autonomous System (AS). Compare with *intra-area route*. See also *AS*, *OSPF*, *route*.

Facilities Data Link—See *FDL*.

FDL—Facilities Data Link. An FDL is a 4-Kbps digital link between a sender and the telephone company's monitors. The link uses Extended Superframe (ESF) framing. The telephone company uses an FDL to check on the quality and performance of T1 lines. You cannot use FDL reporting on a line configured for D4 framing. However, you can obtain D4 and ESF performance statistics in the FDL Stats windows of the Pipeline 220 or in the DSX MIB. See also *ESF*, *T1 line*.

FECN—Forward Explicit Congestion Notification. FECN is a bit set in a Frame Relay header to notify a destination node that there is traffic congestion on the network. Compare with *BECN*. See also *Frame Relay*.

File Transfer Protocol—See *FTP*.

filter—A set of rules describing what action the Pipeline 220 should take when it encounters certain types of packets. A filter can apply to incoming packets, outgoing packets, or both. A packet filter applies to packets on an interface. A route filter applies to routes in Routing Information Protocol (RIP) update packets. See also *packet filter*, *route filter*.

Filter profile—A profile containing parameters that set up filter rules. See also *filter*, *packet filter*, *route filter*.

firewall—See *Secure Access Firewall*.

flash card—See *PCMCIA card*.

flash memory—See *NVRAM*.

foreign agent—An Ascend unit that a mobile node dials into. The foreign agent is the starting point of the Ascend Tunnel Management Protocol (ATMP) tunnel. The foreign agent

must be able to bring up an IP connection to the home agent, and it must authenticate the mobile node by means of a RADIUS user profile. See also *ATMP*, *home agent*.

Forward Explicit Congestion Notification—See *FECN*.

fractional T1 line—A T1 or ISDN BRI line that contains both switched and nailed-up channels. See also *nailed-up channel*, *T1 line*.

frame—In Token Ring, Systems Network Architecture (SNA), and X.25, a packet at the Data Link layer of the OSI Reference Model; in Frame Relay, a packet of fixed size; in Time Division Multiplexing (TDM), a sequence of time slots, each containing a portion of a multiplexed channel. A frame contains source and destination information, flags that designate the start and end of the frame, and information about the integrity of the frame. All other data, such as network protocol information and the actual payload of data, is first encapsulated in a packet. The system then encapsulates the packet in a frame. See also *Data Link layer*, *Frame Relay*, *OSI Reference Model*, *packet*, *TDM*.

framed protocol—A synchronous protocol that encapsulates data into frames. See also *framing*, *protocol*, *synchronous transmission*.

Frame Relay—A WAN architecture originally developed for ISDN lines. A Frame Relay network provides high throughput by handing monitoring functions to higher-level protocols. It is a very efficient standard, with a bandwidth of up to 2 Mbps. Frame Relay is ideal for situations in which periods of very high traffic are interspersed with idle periods. It is protocol independent, and performs routing over virtual circuits called Data Link Connection Indicators (DLCIs). See also *DLCI*.

Frame Relay concentrator—A device that concentrates many low-speed, dial-in connections into one high-speed, nailed-up connection to a Frame Relay switch. When you configure the Pipeline 220 as a Frame Relay concentrator, it accepts incoming dial-in connections as usual and forwards them to a Frame Relay switch. The Pipeline 220 must appear as a Frame Relay switch to both Pipeline 220 users and other Frame Relay switches (such as those from Cascade or Stratacom). See also *Frame Relay switch*.

Frame Relay connection—A link between a dial-in user and the Frame Relay switch. The Pipeline 220 supports the following types of interfaces to the Frame Relay network:

- User-To-Network Interface–Data-Circuit-Terminating-Equipment (UNI-DCE)
- User-To-Network Interface–Data Terminal Equipment (UNI-DTE)

See also *Frame Relay*, *Frame Relay switch*, *UNI-DCE interface*, *UNI-DTE interface*.

Frame Relay profile—A profile that defines the logical link between the Pipeline 220 and a Frame Relay switch. See also *Frame Relay*, *Frame Relay switch*.

Frame Relay switch—A device that sends Frame Relay data out to the Frame Relay network. See also *Frame Relay*.

framing—At the Physical and Data Link layers of the OSI model, a method of fitting bits into a unit called a *frame*. A frame contains source and destination information, flags that designate the start and end of the frame, and information about the integrity of the frame. All other data, such as network protocol information and the actual payload of data, is first encapsulated in a packet. The system then encapsulates the packet in a frame. See also *Data Link layer*, *encapsulation*, *OSI Reference Model*, *packet*, *Physical layer*.

FTP—File Transfer Protocol. FTP is an Application-layer protocol that enables you to transfer files from one device to another over a network. See also *Application layer*.

full duplex—A type of communications configuration in which data can be transmitted in both directions at the same time. Compare with *half duplex*.

gateway—A device or program that provides mapping at all seven layers of the OSI model and translates between two otherwise incompatible networks or network segments. A gateway performs code and protocol conversion to facilitate traffic between data highways of differing architectures. See also *OSI Reference Model*.

gateway connection—A bridging or routing link in a Frame Relay configuration. In a gateway connection, the Pipeline 220 receives an incoming Point-To-Point Protocol (PPP) call, examines the destination IP address, and brings up the appropriate Connection profile to the destination. If the profile specifies Frame Relay encapsulation, a Frame Relay profile, and a Data Link Connection Indicator (DLCI), the Pipeline 220 encapsulates the packets in Frame Relay (RFC 1490) and forwards the data stream out to the Frame Relay switch. The Frame Relay switch uses the DLCI to route the frames. Compare with *circuit connection*, *redirect connection*. See also *Connection profile*, *DLCI*, *Frame Relay*, *Frame Relay connection*, *Frame Relay profile*, *Frame Relay switch*.

gateway mode—An Ascend Tunnel Management Protocol (ATMP) configuration in which the home agent tunnels packets from the foreign agent to the home network across an open WAN connection. The WAN connection must be online. The home agent does not bring up a WAN connection to the home network in response to a packet it receives through the tunnel. For this reason, the home agent must have a nailed-up WAN connection to the home network. Compare with *router mode*. See also *ATMP*, *foreign agent*, *home agent*, *home network*, *nailed-up circuit*.

generic filter—A packet filter that examines the byte- or bit-level contents of a packet and compares them with a value defined in the filter. To use a generic filter effectively, you need to know the contents of certain bytes in the packets you wish to filter. Protocol specifications are usually the best source of such information. Compare with *IP filter*. See also *data filter*, *packet filter*.

Generic Routing Encapsulation—See *GRE*.

GGP—Gateway-To-Gateway Protocol. GGP is a TCP/IP protocol that transfers routing information between gateways. See also *gateway*, *TCP/IP*.

GMT—Greenwich Mean Time. This term has been changed to *Universal Time Configuration (UTC)*. See *UTC*.

GRE—Generic Routing Encapsulation. GRE provides a simple, general-purpose mechanism for encapsulating an arbitrary Network-layer protocol in another arbitrary Network-layer protocol. When a system needs to route data, it first encapsulates the information in a GRE packet. The system then encapsulates the GRE packet in a protocol supported by the network and forwards the packet to its destination.

Greenwich Mean Time—This term has been changed to *Universal Time Configuration (UTC)*. See *UTC*.

half duplex—A type of communications configuration in which data can be transmitted in only one direction at a time. Compare with *full duplex*.

hardware address—An address assigned by the hardware manufacturer and unique to a device.

hardware interface—A hardware link between two devices. A hardware interface has electrical, physical, and functional specifications that determine how two devices communicate. An electrical specification defines the characteristics of the electrical signals. A physical specification might define the number of pins and wires required, and the order in which the pins and wires are laid out. The functional specification instructs the hardware on how to interpret the electrical signals. Examples of commonly used hardware interfaces are RS-232 and V.24. See also *interface*, *RS-232*, *V.24*.

HDLC—High-Level Data Link Control. HDLC is a synchronous, bit-oriented Data Link layer protocol for data transmission. Frame Relay is an example of an HDLC-based packet protocol. HDLC offers half- or full-duplex communications over circuit- or packet-switched networks, allows point-to-point and multipoint configurations, and provides transmission over both wires and wireless media. See also *Data Link layer*, *Frame Relay*, *full duplex*, *half duplex*, *HDLC channel*, *packet switching*, *point-to-point link*.

HDLC channel—High-Level Data Link Control channel. Vital to call routing on the Pipeline 220, HDLC processing removes encapsulation from high-speed incoming data calls, such as those from ISDN Terminal Adapters (TAs). After removing the link's encapsulation, the HDLC channel passes the data stream to the bridge/router. One 192-channel HDLC card supports all switched channels on an eight-port T1 card. See also *HDLC*.

High-Level Data Link Control—See *HDLC*.

home agent—An Ascend unit that represents the terminating part of the Ascend Tunnel Management Protocol (ATMP) tunnel. It must be able to communicate with the home network directly, through another router, or across a nailed-up WAN connection. See also *ATMP*, *home network*, *nailed-up circuit*, *router*.

home network—A private corporate network in an Ascend Tunnel Management Protocol (ATMP) configuration. A private network is one that cannot communicate directly on the Internet. It might be an IPX network, or an IP network with an unregistered network number. See also *ATMP*, *IP network*, *IPX network*, *network number*.

hop—A single message or packet transmission between host and a router, or between two routers. See also *hop count*, *host*, *router*.

hop count—The number of routers through which a packet passes to get from its source to its destination. See also *hop*, *host*, *router*.

host—A computer on a network, also called a *node* or a *station*.

host number—The portion of an IP address that denotes an individual node on a network. The class of an IP address determines which portion of the address belongs to the network number and which portion belongs to the host number. See also *IP address*, *network number*.

host port—A High-Level Data Link Control (HDLC) channel or a digital modem on the Pipeline 220. The Pipeline 220 routes each call it receives to the appropriate host port. See also *digital signal*, *HDLC channel*.

host route—An IP address with a subnet mask of 255.255.255.255, representing a single host rather than a remote router. See also *host*, *route*, *router*, *subnet mask*.

hub—A device that serves as a termination point for multiple hosts, sending signals onto the proper paths. Typically, a hub contains four to eight connectors. In addition to providing connectors for hosts, many hubs include connectors that you can use to link one hub to another.

hybrid LAN—A network in which some links are capable of sending and receiving analog signals, while others handle digital signals. See also *digital signal*.

ICMP—Internet Control Message Protocol. ICMP is an error-reporting mechanism integral to the TCP/IP protocol suite. Gateways and hosts use ICMP to send reports of datagram problems to the sender. ICMP also includes an echo request/reply function that tests whether a destination is available and responding. See also *gateway*, *host*, *TCP/IP*.

IDRP—Inter-Domain Routing Protocol. IDRP is an International Standards Organization (ISO) protocol for routing packets between disparate administrative domains. It is based on the Border Gateway Protocol (BGP). See also *BGP*, *ISO*.

IEC—Interexchange Carrier. An IEC is a type of telephone service that provides long-distance links between local telephone companies. Well-known IECs include AT&T, MCI, and Sprint. Compare with *LEC*.

IEEE—Institute of Electrical and Electronics Engineers. The IEEE is an organization that maintains the standards for 10BaseT and other communications specifications. See also *10Base-T*.

IGMP—Internet Group Management Protocol. IGMP is a protocol implemented by multicast clients and routers. The Pipeline 220 responds as a client to IGMP packets it receives from a Multicast Backbone (MBONE) router. The packets may use IGMP version-1, IGMP version-2, or IGMP Multicast Trace (MTRACE). Pipeline 220 clients wanting MBONE service must implement IGMP. See also *MBONE*, *multicast*, *multicast network*, *router*.

IGP—Interior Gateway Protocol. IGP transmits routing information internal to a network. See also *routing*.

index—A name, physical address, or interface address that identifies a specific profile of a particular type.

input filter—A filter applied to an incoming packet. See also *filter*, *packet filter*, *route filter*.

Institute of Electrical and Electronics Engineers—See *IEEE*.

Inter-Domain Routing Protocol—See *IDRP*.

Interexchange Carrier—See *IEC*.

interface—A connection between two devices, programs, or program elements. See also *hardware interface*.

interface-based routing—An IP-routing method in which each physical or logical interface on the unit has its own IP address. The interface becomes a numbered interface. Reasons for using numbered interfaces include troubleshooting nailed-up 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 220 to operate more as a multi-homed host behaves.

You can configure each link as numbered (interface-based) or unnumbered (system-based). If no interfaces are numbered, the Pipeline 220 operates as a purely system-based router. Compare with *system-based routing*, *unnumbered interface*. See also *IP routing*, *multi-homed host*, *numbered interface*, *point-to-point link*.

interface table—A table containing the addresses of each Ethernet, IP, and IPX interface on the Pipeline 220. Each packet-handling slot card operates as a router subsystem with its own local interface table and route cache. To view the interface table, enter the Netstat command with the **-in** argument, as shown in the following example:

```
admin> netstat -in
```

Name	MTU	Net/Dest	Address	Ipkts	Ierr	Opkts	Oerr
ie0	1500	12.65.212.0/24	12.65.212.227	107219	0	54351	0
lo0	1500	127.0.0.1/32	127.0.0.1	4867	0	4867	0
rj0	1500	127.0.0.2/32	127.0.0.2	0	0	0	0
bh0	1500	127.0.0.3/32	127.0.0.3	0	0	0	0
wan4	1500	10.122.99.1	-	0	0	0	0
ie1-12-1	1500	11.168.6.0/24	11.168.6.227	430276	651	0	0
ie1-12-2	1500	10.122.72.0/24	10.122.72.1	0	0	0	3144
ie1-12-3	1500	10.122.73.0/24	10.122.73.1	0	0	3142	0
ie1-12-4	1500	10.122.74.0/24	10.122.74.1	0	0	3141	0

The following table describes each column in the interface table.

Column	Indicates
Name	Internal name of the interface.
MTU	Maximum Transfer Unit, the largest packet that can be transmitted over the interface. If a packet's size exceeds the MTU, the packet must be fragmented or segmented, and then reassembled at the receiving end.
Net/Dest	IP address of the destination of packets on the interface.
Address	IP address of the interface.
Ipkts	Number of incoming packets.
Ierr	Number of errors recorded for incoming packets.
Opkts	Number of outgoing packets.
Oerr	Number of errors recorded for outgoing packets.

The following table describes the types of interfaces that can appear in the interface table.

Entry	Indicates
$ie0$ or $ieN-N-N-N[-N]$	Ethernet interfaces, where $N-N-N-N$ represents the shelf-number, slot-number, item number, and logical item number of the interface. When the logical item number is zero, it does not appear in the interface name.
$lo0$	Loopback interface.
$rj0$	Reject interface used for the pool summary feature.
$bh0$	Black-hole interface used for the pool summary feature
$wanN$	WAN connection, entered as the connection becomes active.

See also *POP*.

Interior Gateway Protocol—See *IGP*.

International Standards Organization—See *ISO*.

International Telecommunication Union—Telecommunication Standardization

Sector—See *ITU-T*.

internet—A series of networks connected by bridges, gateways, or routers. An internet is also called an *internetwork*. See also *bridge*, *gateway*, *router*.

Internet—The complex of WANs joining government, university, corporate and private computers in a vast web of network interconnection.

Internet Control Message Protocol—See *ICMP*.

Internet gateway—A gateway for accessing the Internet. See also *gateway*.

Internet Group Management Protocol—See *IGMP*.

Internet Network Information Center—See *InterNIC*.

Internet Protocol—See *IP*.

Internet Protocol Control Protocol—See *IPCP*.

Internet Service Provider—See *ISP*.

internetwork—See *internet*.

Internetwork Packet Exchange—See *IPX*.

InterNIC—Internet Network Information Center. InterNIC is an organization that provides Internet information services, oversees the registration of Internet addresses and Domain Name System (DNS) names, assigns RFC numbers, and assists users in gaining access to the Internet. See also *DNS*, *RFC*.

interoperability—Compatibility with the devices and services of multiple vendors. Interoperable devices can be integrated into a network containing a wide range of vendor products. Interoperability is a significant factor among expansion considerations, because any device must have the versatility to function in an expanding network structure. The technical elements of interoperability may include a bundle of protocols and a flexible architecture to accommodate upgrades. A remote-access server should include capabilities such as translation, encapsulation, and filtering.

intra-area route—A route imported into the Open Shortest Path First (OSPF) database from within the router's area. Compare with *external route*. See also *area*, *OSPF*, *route*, *router*.

IP—Internet Protocol. IP provides connectionless, non-guaranteed transmission of Transport-layer data packets. IP fragments packets, allowing them to take different paths across the WAN, and then reassembles them into the proper order at their destination. See also *Transport layer*.

IP address—An address that uniquely identifies each host on a network or internet. An IP address has a length of 32 bits, and is divided into four 8-bit parts, each separated by a period, as in 149.122.3.30. This kind of notation is called *dotted decimal notation*. Each part can consist of a number between 1 and 255.

An IP address consists of a network number and a host number. IP addresses come in three types: Class A, Class B, and Class C. The class of an IP address determines which portion of the address belongs to the network number and which portion belongs to the host number. The first bits of the IP address identify the class. The Internet Network Information Center (InterNIC) determines the type of class assigned a network.

A Class A address starts with 0 as the class identifier, followed by 7 bits for the network number and 24 bits for the host number. Therefore, the first number in dotted decimal form is the network number. The next three numbers make up the host number. For example, in the IP address 127.120.3.8, the network number is 127 and the host number is 120.3.8. This type of address is used by the largest organizations, because this scheme allows for over 16 million different host numbers. However, it also limits network numbers to a total of 128.

A Class B address starts with binary 10 as the class identifier, followed by 14 bits for the network number and 16 bits for the host number. Therefore, the first two dotted decimal numbers comprise the network number, and the second two dotted decimal numbers comprise the host number. For example, in the IP address 147.14.86.24, the network number is 147.14 and the host number is 86.24. More network numbers are available than in a Class C address, but fewer hosts (approximately 65,000).

A Class C address starts with binary 110 as the class identifier, followed by 21 bits for the network number and 9 bits for the host number. Therefore, the first three dotted decimal numbers comprise the network number, and the last dotted decimal number comprises the host number. For example, in the IP address 225.135.38.42, the network number is 225.135.38 and the host number is 42. Many network numbers are available, but only 254 hosts per network number. The numbers 0 and 255 are reserved.

You can tell the type of class an IP address falls into by looking at the first 8-bit portion of the dotted decimal form of the address. Class A addresses begin with a number between 0 and 127. Class B addresses begin with a number between 128 and 223. Class C addresses begin with a number between 192 and 233.

In addition to an IP address, you can use a symbolic name provided by Domain Name System (DNS) to designate an Internet address.

See also *DNS*, *dotted decimal notation*, *host number*, *internet*, *InterNIC*, *IP*, *network*, *network number*.

IP address spoofing—A way for a remote device to illegally acquire a local address in order to break through a firewall or data filter.

IPCP—Internet Protocol Control Protocol. IPCP is a protocol for configuring, enabling, and disabling the IP protocol modules on both ends of a point-to-point link. IPCP is tied to PPP, and is activated only when PPP reaches the Network-layer protocol phase. IPCP packets received prior to this phase are discarded. Elements of IPCP include packet encapsulation, code fields, and timeouts. See also *IP*, *Network layer*, *point-to-point link*.

IP filter—A packet filter that examines fields specific to IP packets. An IP filter focuses on known fields, such as source or destination address and protocol number. It operates on logical information that is relatively easy to obtain. In an IP filter, a number of distinct comparisons occur in a defined order. When a comparison fails, the packet goes on to the next comparison. When a comparison succeeds, the filtering process stops and the Pipeline 220 applies the forward action in the filter to the packet. Compare with *generic filter*. See also *data filter*, *packet filter*.

IP multicast forwarding—See *multicast forwarding*.

IP network—A network that uses the Internet Protocol (IP) to transmit packets at the Transport layer. See also *IP*.

IP redirection—A configuration in which the Pipeline 220 automatically redirects incoming IP packets to a host you specify on the local IP network. When you specify IP redirection, the Pipeline 220 bypasses all internal routing and bridging tables, and sends all packets it receives on a connection's WAN interface to the specified IP address. IP redirection does not affect outgoing packets. See also *bridging table*, *IP address*, *IP routing table*.

IP route—A path from one IP network to another. See also *dynamic route*, *IP network*, *multipath route*, *static route*.

IP router—A device that sends IP packets from a source to a destination by multiple paths. As an IP router, the Pipeline 220 routes IP packets between its Ethernet interfaces and across any WAN interface configured for IP routing. See also *IP route*, *IP routing*.

IP routing—A method of determining how to forward an IP packet to the proper destination. When acting as an IP router, the Pipeline 220 routes IP packets between its Ethernet interfaces and across any WAN interface configured for IP routing. It consults its internal routing table to determine where to forward each IP packet it processes. First, the Pipeline 220 tries to find a match between the packet's destination address and a Destination field in its routing table. If it finds a match, it looks brings up the required connection (if necessary) to reach the next-hop router specified for that route, and forwards the packet.

If it does not find a match for the packet's destination address, it looks for a default route (destination address 0.0.0.0). If it finds a default route, it brings up the required connection (if necessary) and forwards the packet. If the routing table has no default route, and no route that matches a packet's destination address, the Pipeline 220 drops the packet.

See also *default route*, *hop*, *IP route*, *IP router*, *IP routing table*.

IP routing table—A table that contains information about how to forward IP packets. On the Pipeline 220, the routing table contains the fields Destination, Gateway, IF, Flg, Pref, Metric, Use, and Age. on the Pipeline 220, the routing table contains the following fields:

Field	Indicates
Destination	Target address. To send a packet to this address, the Pipeline 220 uses the route. Note that the router uses the most specific route (having the largest subnet mask) that matches a given destination.
Gateway	Address of the next-hop router that can forward packets to the destination. Direct routes do not show a gateway address.
IF	Name of the interface through which the Pipeline 220 sends a packet addressed to the destination.
Flg	Flag values describing the route: <ul style="list-style-type: none"> • C (A directly connected route, such as Ethernet) • I (An ICMP Redirect dynamic route) • N (A route placed in the table by the SNMP MIB II) • O (A route learned from OSPF) • R (A route learned from RIP) • r (A transient RADIUS-like route that will disappear when the connection drops) • S (A static route) • ? (A route of unknown origin, which indicates an error) • G (An indirect route through a gateway) • P (A private route) • T (A temporary route) • M (A multipath route) • * (A backup static route for a transient RADIUS-like route)
Pref	Preference value of the route.
Metric	RIP-style metric for the route, with a valid range of 0-16. Routes learned from OSPF show a RIP metric of 10. OSPF Cost infinity routes show a RIP metric of 16.
Use	Count of the number of times the route was referenced since it was created. (Many of these references are internal, so this is not a count of the number of packets sent using the route.)
Age	Age of the route in seconds. It is used for troubleshooting, and to determine when routes are changing rapidly.

See also *direct route*, *dynamic route*, *gateway*, *hop*, *IP route*, *IP router*, *metric*, *multipath route*, *OSPF*, *preference*, *RIP*, *static route*.

IP subnet—A portion of an IP network. IP subnetting is a way to subdivide a network into smaller networks, resulting in a greater number of hosts on a network associated a single IP network number. An IP address that uses a subnet has three elements: network, subnet, and host. You identify a subnet by combining an address with a subnet mask. For example, in the address 195.112.56.75/14, /14 is the subnet mask. See also *host number*, *IP address*, *network number*, *subnet mask*.

IP switch—A device that can determine the destination of large volumes of incoming IP packets and send them to the appropriate outgoing ports at high speeds. An IP switch is a high-performance device designed for high-volume, large-scale public and private backbone applications. See also *switch*.

IPX—Internetwork Packet Exchange. IPX is Novell's connectionless Network-layer protocol. Derived from XNS' Internetwork Datagram Protocol (IDP), IPX performs addressing and routing functions. At the server, IPX passes outgoing datagrams to the network interface software. At the packet's destination, IPX passes the data to upper-layer processes. Along an IPX route, intermediate devices use IPX to route packets to their destinations. When routing, IPX relies on information supplied by the Routing Information Protocol (RIP). See also *IPX network*, *IPX route*, *IPX routing*, *IPX server*, *RIP*.

IPX bridging—At the Data Link layer, a way of passing IPX packets between networks. See also *Data Link layer*, *IPX network*.

IPX client—A user or device who gains access to the services of an IPX server. See also *IPX server*.

IPXCP—Internet Packet Exchange Control Protocol. IPXCP is a protocol for configuring, enabling, and disabling the IPX protocol modules on both ends of a point-to-point link. IPXCP is tied to PPP, and is activated only when PPP reaches the Network-layer protocol phase. IPXCP packets received prior to this phase are discarded. Elements of IPXCP include packet encapsulation, code fields, and timeouts. See also *IPX*, *point-to-point link*.

IPX frame—The type of packet frame used by an IPX server. An IPX frame can follow the IEEE 802.2, IEEE 802.3, SubNetwork Access Protocol (SNAP), or Ethernet II protocol specification for the Media Access Control (MAC) header. See also *802.2*, *802.3*, *Ethernet II*, *IPX server*, *MAC*, *SNAP*.

IPX network—A network consisting of one or more IPX servers and IPX clients. See also *IPX client*, *IPX server*.

IPX route—A path from one IPX network to another. See also *IPX network*, *IPX router*.

IPX router—A device that sends IPX packets from a source to a destination by various paths. See also *IPX route*.

IPX routing—A method of sending IPX packets from a source to a destination at the Network layer. See also *Network layer*.

IPX server—A server that runs the NetWare operating system, manages network resources, and communicates with IPX clients. See also *IPX*, *IPX client*.

island—A group of networks on the Multicast Backbone (MBONE). The islands are connected by tunnels and support IP. See also *MBONE*.

ISO—International Standards Organization. The ISO is an organization devoted to the definition of standards for national and international data communications. The U.S. representative to the ISO is the American National Standards Institute (ANSI). Companies whose products are ISO certified reflect a high quality of consistency and quality.

ISO 9001—The current set of International Standards Organization (ISO) standards. See also *ISO*.

ISP—Internet Service Provider. An ISP is a company that provides access to the Internet. By establishing Points of Presence (POPs) containing remote-access servers and a suite of user software packages, the ISP acts as a commercial on-ramp to the Internet. Providers typically charge a monthly fee, and supply technical support and advice to customers.

ITU-T—International Telecommunication Union–Telecommunication Standardization Sector. The ITU-T is the committee that replaced the Consultative Committee for International Telegraphy and Telephony (CCITT) on March 1, 1993. The ITU-T is responsible for a wide array of telecommunications and networking standards.

Java—An object-oriented programming language developed by Sun Microsystems, Inc. You can use Java to create applets for distribution on the World Wide Web. Java programs run inside a Java-enabled Web browser or inside a Java Virtual Machine (JVM).

Java Virtual Machine—See *JVM*.

JVM—Java Virtual Machine. A JVM is an abstract computer that runs compiled Java code. The JVM is “virtual” because it is software that runs on top of a hardware platform and an operating system. All Java programs are compiled for a JVM. See also *Java*.

LAN—Local Area Network. A LAN is a network in which two or more computers, located within a limited distance of one another, are connected in order to share files and resources. A PC-based LAN consists of a dedicated server running a network operating system and attached to several workstations. A host-based LAN consists of one or more hosts and terminals. Examples of LAN architectures are Ethernet, ARCnet, Fiber Distributed Data Interface (FDDI), and Token Ring.

LAN adapter—See *NIC*.

LAN packet display—A display of packet performance over a specified time, measured graphically or by counters.

LAN/WAN connectivity—The ability to link Local Area Networks (LANs) and Wide Area Networks (WANs). A wide range of tools, from translation protocols to communications features to support services, make a remote-access device like the Pipeline 220 an effective link between LANs and WANs.

LAP—Link Access Procedure. LAP is a protocol containing a subset of High-Level Data Link Protocol (HDLC) features. In order to maintain compatibility with HDLC, LAP was changed to create LAPB. See also *LAPB*.

LAPB—Link Access Procedure, Balanced. LAPB is a protocol for B channels that use packet-switching mode. See also *packet switching*.

LAPD—Link Access Procedure, D channel. LAPD is a protocol for the D channel. It provides the mechanism for combining multiple channels into a single logical link, and for monitoring and controlling the flow of data over the B channels. See also *DDP*.

LAPF—Link Access Procedure, Frame. LAPF is a protocol for Frame-mode bearer services.

LAPM—Link Access Procedure, Modem. LAPM is an error-detection protocol for correcting data communication errors occurring on the link between two modems.

latency—For a communications channel, the amount of time before the channel is available for a transmission; for data transmissions, the amount of time it takes for a packet to reach its destination. The following elements contribute to latency:

- The type of physical media in use.
- Physical interference from noise or other signals.
- Required setup and teardown times.
- Signal interfaces. Ethernet consumes a minimum of 0.3 milliseconds (ms). A 28.8 modem takes about 300 times longer.
- Bottlenecks, such as the 50 ms it takes to move data through a serial port.
- Data conversion, such as the conversion from digital to analog data required by a modem.
- Compression.

Once latency is present, it cannot be optimized. You must remove the cause. To maximize throughput, use the highest bandwidth available. All services go as fast as the medium allows. For example, if the medium is copper, the speed of the electrical signal through the copper does not vary with the type of line in use. A T1 line is considered faster than a single analog line only because its bandwidth is greater.

LCP—Link Control Protocol. LCP sets up, manages, and tears down a connection between two Point-To-Point Protocol (PPP) endpoints. See also *PPP*.

learning bridge—See *transparent bridge*.

leased circuit—See *nailed-up circuit*.

leased line—See *nailed-up line*.

LEC—Local Exchange Carrier. An LEC is a local telephone company. See also *IEC*.

line—A physical interface to the WAN. A line consists of one or more channels, each of which can transmit data. See also *channel*.

Line Quality Monitoring—See *LQM*.

Link Access Procedure—See *LAP*.

Link Access Procedure, Balanced—See *LAPB*.

Link Access Procedure, D Channel—See *LAPD*.

Link Access Procedure, Frame—See *LAPF*.

Link Access Procedure, Modem—See *LAPM*.

link compression—A process that removes waste and redundancy from the data on a connection, enabling faster throughput. For the Pipeline 220 to use link compression, both sides must be configured to use the same compression method. You can use Stac compression (an Ascend-modified version of draft 0 of the CCP protocol), Stac-9 compression (the method

specified by draft 9 of the Stac LZS compression protocol), or Microsoft Stac compression (the method implemented by Windows 95). See also *CCP*, *slot compression*, *VJ compression*.

Link Control Protocol—See *LCP*.


link state—The condition of an Open Shortest Path First (OSPF) link. See also *OSPF*.

Link-State Advertisement—See *LSA*.

link-state database—A database that contains Open Shortest Path First (OSPF) routing information. Link-state routing algorithms require that all routers within a domain maintain identical link-state databases, and that the databases describe the complete topology of the domain. An OSPF router's domain may be an Autonomous System (AS) or an area within an AS.

Based on the exchange of information among routers, OSPF routers create a link-state database, which is updated based on packet exchanges among the routers. Link-state databases are synchronized between pairs of adjacent routers. In addition, each OSPF router uses its link-state database to calculate a self-rooted tree of shortest paths to all destinations. The routing table is built from these calculated shortest-path trees. Externally derived routing data is advertised throughout the AS but is kept separate from the link-state data. Each external route can also be tagged by the advertising router, enabling the passing of additional information between routers on the boundary of the AS.

For example, suppose you have the network topology in the following illustration:



The link-state databases of the three routers contain the cost information, as follows:

Router-1

Network-1/Cost 0

Network-2/Cost 0

Router-2/Cost 20

Router-2

Network-2/Cost0

Network-3/Cost0

Router-1/Cost 20

Router-3/Cost 30

Router-3

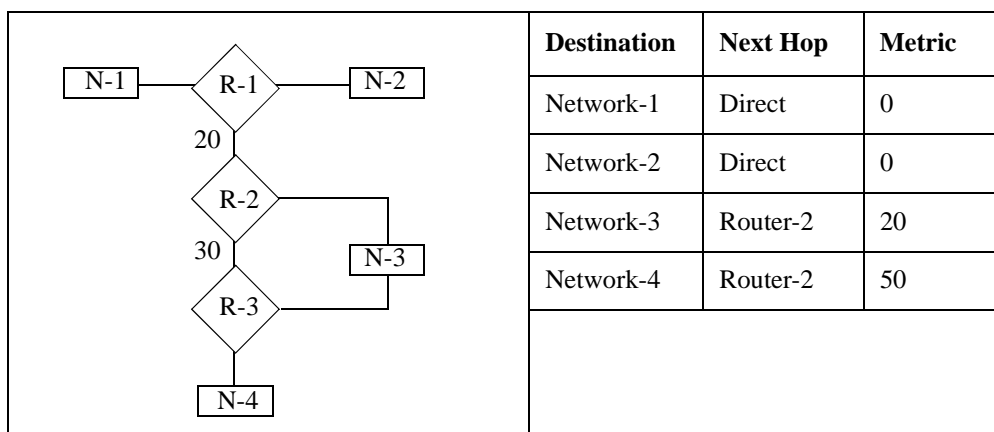
Network-3/Cost 0

Network-4/Cost 0

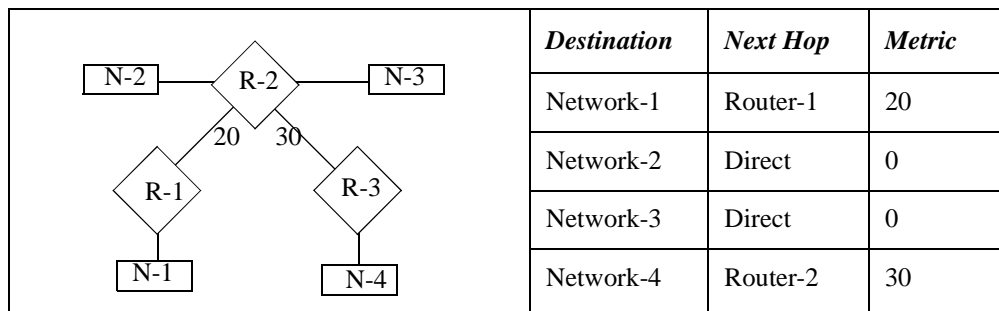
Router-2/Cost 30

Each router builds a self-rooted shortest-path tree, and then calculates a routing table stating the shortest path to each destination in the AS.

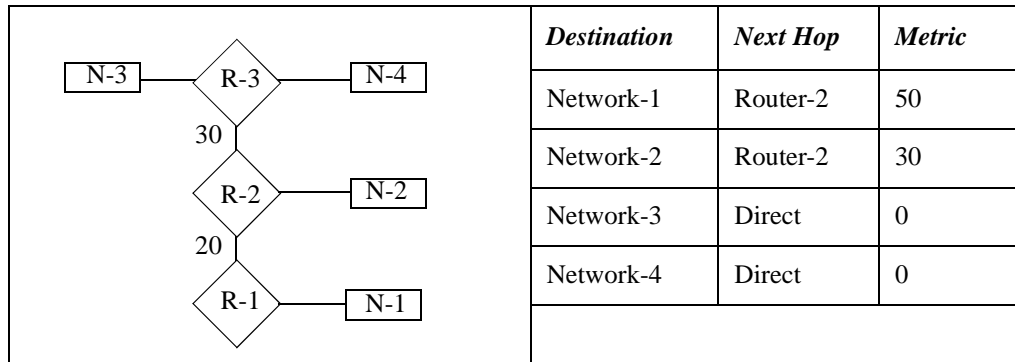
The following figure displays the shortest-path tree and resulting routing table for Router-1:



The following figure displays the shortest-path tree and resulting routing table for Router-2:



The following figure displays the shortest-path tree and resulting routing table for Router-3:



See also *adjacency*, *AS*, *OSPF*.

link-state metric—A metric that takes into account a variety of link conditions, such as the reliability or speed of the link, when determining the best path to a destination network. Open Shortest Path First (OSPF) is a link-state protocol. Compare with *distance-vector metric*. See also *OSPF*.

Link-State-Request packet—An Open Shortest Path First (OSPF) request for an updated database. To make routing decisions, OSPF uses a link-state database of the network and propagates only changes to the database. See also *link-state database*, *OSPF*, *routing*.

Link-State-Update packet—A packet exchanged between Open Shortest Path First (OSPF) routers for the purpose of updating link-state databases. See also *OSPF*, *router*.

LLC—Logical Link Control. In the IEEE's Local Area Network/Reference Model, LLC denotes a sublayer above the Media Access Control (MAC) sublayer. Combined, the LLC and MAC sublayers are equivalent to the Data Link layer in the OSI Reference Model. They give higher-level protocols access to the physical media. See also *MAC*, *OSI Reference Model*.

Local Exchange Carrier—See *LEC*.

Local Area Network—See *LAN*.

local device—A device directly connected to the Ascend unit or residing on the local Ethernet.

logical address—An address assigned by a network administrator to associate several devices with one another into a logical hierarchy or group. A router uses the logical address to help transmit a packet to its destination. An example of a logical address is an IP address. Compare with *hardware address*. See also *IP address*, *router*.

logical item number—In an interface address, the number identifying a specific logical interface or channel on a physical line or port. The logical item number is zero except when the device has multiple interfaces or supports multiple channels. For example, a T1 line may support 24 channels, each of which is specified in a Call-Route profile by an interface address containing a logical item number from 1 to 24. See also *T1 line*.

logical link—The link between the Pipeline 220 and a Frame Relay switch, as defined in a Frame Relay profile. See also *Frame Relay*, *Frame Relay profile*, *Frame Relay switch*.

Logical Link Control—See *LLC*.

log level—The level of event information the Pipeline 220 displays at the console.

LQM—Line Quality Monitoring. LQM is a feature that enables the Pipeline 220 to monitor the quality of a link. When you enable LQM, the Pipeline 220 counts the number of packets sent across the link and periodically asks the remote end how many packets it has received. Discrepancies are evidence of packet loss and indicate link-quality problems. The Pipeline 220 can tear down and reestablish a call if the problems on the link exceed a specified threshold.

LSA—Link-State Advertisement. An LSA is a packet that describes various aspects of an Open Shortest Path First (OSPF) route. OSPF defines the following LSA types:

LSA Type	Description
Type 1 (RTR)	Router-LSA that describes the collected states of the router's interfaces.
Type 2 (NET)	Network-LSA that describes the set of routers attached to the network.
Types 3 and 4 (STUB)	Summary-LSA that describes point-to-point routes to networks or Area Border Routers (ABRs).
Type 5 (ASE)	AS-external-LSA that describes routes to destinations external to the AS. An AS-external-LSA can also describe a default route for the AS.

See also *AS*, *ASE*, *ASE Type-5*, *OSPF*, *point-to-point link*, *route*, *router*.

MAC—Media Access Control. In the IEEE's Local Area Network/Reference Model, MAC denotes a sublayer below the Logical Link Control (LLC) sublayer. Combined, the LLC and MAC sublayers are equivalent to the Data Link layer in the OSI Reference Model. They give higher-level protocols access to the physical media. See also *LLC*, *MAC address*, *OSI Reference Model*.

MAC address—The 6-byte hexadecimal address that the manufacturer assigns to the Ethernet controller for a port. See also *hardware address*, *MAC*.

Management Information Base—See *MIB*.

manager—An application that receives Simple Network Management Protocol (SNMP) information from an agent. An agent and manager share a database of information, called the Management Information Base (MIB). An agent can use a message called a traps-PDU to send unsolicited information to the manager. A manager that uses the Ascend Enterprise MIB can query the Pipeline 220, set parameters, sound alarms when certain conditions appear, and perform other administrative tasks. See also *agent*, *community name*, *MIB*, *SNMP*, *traps-PDU*.

mask—In a generic filter, a 12-byte value the Pipeline 220 applies to a packet before comparing its contents to the value you indicate in a filter specification. The mask hides the bits that appear behind each binary 0 (zero). A mask of all ones (FF:FF:FF:FF:FF:FF:FF:FF)

masks no bits, so the full specified value must match the packet contents. See also *generic filter*.

Maximum Receive Unit—See *MRU*.

Maximum Transfer Unit—See *MTU*.

MBONE—Multicast Backbone. The MBONE is a virtual network layered on top of the Internet to support IP multicast routing across point-to-point links. Because multicasting is a fast and inexpensive way to communicate information to multiple hosts, the MBONE is used for transmitting audio and video on the Internet in real time.

The MBONE consists of groups of networks called *islands*. These islands are connected by tunnels and support IP. When the Pipeline 220 accesses an MBONE network, it starts receiving MBONE multicasts. It resends the multicast packets to all of its own clients connected to it for MBONE service. The clients wanting MBONE service must implement Internet Group Membership Protocol (IGMP).

To the MBONE, the Pipeline 220 looks like a multicast client. It responds as a client to IGMP packets it receives from an MBONE router. The MBONE router can reside on the Pipeline 220 unit's Ethernet interface or across a WAN link. If the router resides across a WAN link, the Pipeline 220 can respond to multicast clients on its Ethernet interface as well as across the WAN.

To multicast clients on a WAN or Ethernet interface, the Pipeline 220 looks like a multicast router, although it simply forwards multicast packets on the basis of group memberships. See also *multicast*, *multicast forwarding*, *multicast heartbeat*, *multicast network*, *multicast rate limit*, *point-to-point link*.

MBONE interface—The location on the Pipeline 220 that connects to an MBONE router. See also *MBONE router*.

MBONE router—A router that directs multicast packets to a group of clients on a subscription list. See also *MBONE*, *MBONE interface*, *multicast*, *multicast forwarding*, *multicast heartbeat*, *multicast network*, *multicast rate limit*.

Media Access Control—See *MAC*.

menu mode—A mode in which the terminal server presents a banner message and a menu of hosts. In menu mode, a user cannot enter terminal-server commands, but can connect by means of Telnet, Rlogin, or raw TCP to the hosts you specify. If you configure the menu locally, you can specify up to four hosts. Compare with *command mode*.

message—Data transmitted from one location to another with a header field, information field, and trailer. Often used interchangeably with *packet* and *frame*.

metric—A value that determines how quickly a packet can reach its destination. Routing Information Protocol (RIP) and the Open Shortest Path First (OSPF) protocol use different types of metrics.

- RIP is a distance-vector protocol, which uses a hop count to select the shortest route to a destination network. RIP always uses the lowest hop count, regardless of the speed or reliability of a link.
- OSPF is a link-state protocol. OSPF can take into account a variety of link conditions, such as the reliability or speed of the link, when determining the best path to a destination network.

See also *hop count*, *OSPF*, *preference*, *RIP*, *route*.

MIB—Management Information Base. A MIB is a Simple Network Management Protocol (SNMP) database of information available to network management programs. An agent creates a MIB. A network manager queries the MIB for information, and might create a MIB of its own. The MIB on the agent contains machine-specific information. The manager's MIB has more general information. The Pipeline 220 supports SNMP MIB II, T1 MIB, and Ascend Enterprise MIBs. See also *agent*, *manager*, *SNMP*.

Microsoft CHAP—See *MS-CHAP*.

Microsoft Stac—The version of the Stac LZS compression method implemented by Windows 95. Compare with *Stac compression*, *Stac-9 compression*.

MRU—Maximum Receive Unit. An MRU is the largest packet that a host on a link can receive. Compare with *MTU*.

MS-CHAP—Microsoft CHAP. MS-CHAP is a close derivative of Challenge Handshake Authentication Protocol (CHAP). However, CHAP is designed to authenticate WAN-aware secure software, and is not widely used to support remote workstations, where an insecure plain text login might be required. MS-CHAP addresses this issue, and also integrates the encryption and hashing algorithms used on Windows networks. Microsoft Windows NT and LAN Manager platforms implement MS-CHAP. Compare with *CHAP*.

MTU—Maximum Transfer Unit. An MTU is the largest packet that can be transmitted over a particular medium. If a packet's size exceeds the MTU, the packet must be fragmented or segmented, and then reassembled at the receiving end. Compare with *MRU*.

multicast—A transmission method in which one device communicates with destination hosts by means of a single transmission to all recipients of a subscriber list. See also *MBONE*, *multicast forwarding*, *multicast heartbeat*, *multicast network*, *multicast rate limit*.

Multicast Backbone—See *MBONE*.

multicast forwarding—A process by which the Pipeline 220 forwards traffic it receives on one of its Ethernet or WAN interfaces from an Multicast Backbone (MBONE) router. To the MBONE, the Pipeline 220 looks like a multicast client, and it responds as a client to Internet Group Membership Protocol (IGMP) packets it receives. The Pipeline 220 resends the multicast packets to all of its own clients connected to it for MBONE service. The clients wanting MBONE service must implement IGMP.

Each Ethernet or WAN interface that supports multicasting must be configured to allow multicasting forwarding. When you do so, the Pipeline 220 begins handling IGMP requests and responses on the interface. It does not begin forwarding multicast traffic until you set the multicast rate limit. See also *IGMP*, *MBONE*, *MBONE router*, *multicast*, *multicast heartbeat*, *multicast network*, *multicast rate limit*.

multicast heartbeat—A feature that enables you to monitor possible connectivity problems. Using the multicast heartbeat feature, you configure the Pipeline 220 to poll continuously for multicast traffic. The Pipeline 220 generates the following SNMP alarm trap if a traffic breakdown occurs:

```
Trap type:  TRAP_ENTERPRISE
Code:      TRAP_MULTICAST_TREE_BROKEN (19)
```

Arguments:

- 1) Multicast group address being monitored (4 bytes),
- 2) Source address of last heartbeat packet received (4 bytes)
- 3) Slot time interval configured in seconds (4 bytes),
- 4) Number of slots configured (4 bytes).
- 5) Total number of heartbeat packets received before the MAX started sending SNMP Alarms (4 bytes).

Heartbeat monitoring is optional. It is not required for multicast forwarding. To set up heartbeat monitoring, you configure several parameters that define what packets will be monitored, how often the Pipeline 220 polls for multicast packets, and what threshold must be reached for the Pipeline 220 to generate an alarm.

See also *MBONE*, *multicast*, *multicast forwarding*, *multicast network*, *multicast rate limit*, *SNMP*.

multicast network—A network in which a router sends packets to all addresses on a subscriber list. This type of network is different from both a unicast network (in which the router sends packets to one user at a time) and a broadcast network (in which the router sends packets to all users, whether they appear on subscription lists or not). The Multicast Backbone (MBONE) is an example of a multicast network. See also *MBONE*, *multicast*, *multicast forwarding*, *multicast heartbeat*, *multicast rate limit*.

multicast rate limit—A way to limit the rate at which the Pipeline 220 accepts multicast packets from its clients. To begin forwarding multicast traffic on the MBONE interface, you must set the multicast rate limit to a number less than 100. For example if you set the limit to 5, the Pipeline 220 accepts a packet from multicast clients on the interface every 5 seconds. Any subsequent packets received in that 5-second window are discarded. See also *MBONE*, *MBONE interface*, *multicast*, *multicast forwarding*, *multicast heartbeat*, *multicast network*.

multi-homed host—A single Internet device connected to multiple data paths. Each link may reside on a different network.

multipath route—A static route that distributes the traffic load across multiple interfaces to a single destination. See also *route*, *static route*.

nailed-up channel—A channel on a line rented from the phone company for exclusive use, 24 hours per day, seven days per week. See also *nailed-up circuit*, *nailed-up line*.

nailed-up circuit—A permanent connection between endpoints over which two parties exchange data. The number of nailed-up channels must be the same at both ends of the connection. For example, if there are five nailed-up channels at the local end, there must be five nailed-up channels at the remote end. However, channel assignments do not have to match. For example, channel 1 may be switched at the local end and nailed up at the remote end. A nailed-up circuit is also known as a *private circuit* or a *leased circuit*. See also *nailed-up channel*, *nailed-up line*.

nailed-up line—A line rented from the phone company for exclusive use, 24 hours per day, seven days per week. The connection exists between two predetermined points and cannot be switched to other locations. A nailed-up line is also called a *leased line*. See also *nailed-up channel*, *nailed-up circuit*.

name and password authentication—A form of authentication in which the Pipeline 220 attempts to match a caller's user name and password to the parameters or attributes

specified in a profile. If name and password authentication is required, the Pipeline 220 first attempts to match the caller's name and password to a local Connection profile. If authentication succeeds using a local Connection profile, the Pipeline 220 uses the parameters specified in the profile to build the connection.

See also *authentication*, *Connection profile*.

NAS—Network Access Server. An NAS is a device that provides LAN and WAN access for network hosts. The Pipeline 220 is an example of an NAS.

NAT for LAN—Network Address Translation for LAN. NAT for LAN is a feature that allows a Pipeline to connect a LAN to a remote network, even if devices on the LAN have addresses that are not valid for the remote network. The Pipeline translates between the local network addresses and the remote network addresses.

Access to public networks requires the use of an official IP address that is unique across the entire network. Typically, a central authority assigns a range of addresses, and a local administrator distributes them. If access to a public network is not necessary, the local manager can assign addresses as he or she sees fit, even if the addresses are unofficial or belong to another company.

Because the supply of addresses is rapidly diminishing, a company might not be able to get official addresses for its entire network. A site might already have unofficial addresses, but now needs access to the Internet, where an official address is required. For these reasons, you might need a facility for borrowing an official address and dynamically translating between the local and official addresses. NAT for LAN provides this facility.

See also *IP address*.

NCP—NetWare Core Protocol. NCP is a protocol that allows an IPX server to respond to client requests. See also *IPX server*.

NCP—Network Control Protocol. NCP is a collection of protocols for setting up and configuring Network-layer protocols (such as AppleTalk) over PPP. See also *PPP*.

NetBIOS—Network Basic Input/Output System. NetBIOS is a protocol developed by IBM that provides network access to upper-layer programs. NetBIOS functionality includes the Session, Presentation, and Application layers of the OSI Reference Model, and provides naming services, connectionless best-effort datagram delivery, and support for virtual circuits. See also *OSI Reference Model*.

NetWare Core Protocol—See *NCP*.

NetWare server—See *IPX server*.

network—A group of computers, often called *hosts*, *nodes*, or *stations*, that are connected to each other for the purpose of sharing files and other resources. Each computer has a Network Interface Card (NIC) that enables it to gain access to the network. Each host can have one or more peripherals (such as a fax modem or printer) attached to it. Each peripheral can be shared with other network users, or can remain private to the individual computer.

Network Access Server—See *NAS*.

network adapter—See *NIC*.

Network Address Translation for LAN—See *NAT for LAN*.

network address—An address shared by all the hosts on the same physical network.

network alignment—A method of setting up IP address pools for pool summary. When you perform network alignment, you make sure that the first address in the pool is the first host address, and that the maximum number of entries you specify is two fewer than the total number of addresses in the pool. See also *IP address*, *POP*.

Network Basic Input/Output System—See *NetBIOS*.

network board—See *NIC*.

Network Control Protocol—See *NCP*.

Network File System—See *NFS*.

Network Information Center—See *InterNIC*.

Network Information Service—See *NIS*.

Network Interface Card—See *NIC*.

Network layer—A layer in the OSI Reference Model. The Network layer provides address resolution and routing protocols. Address resolution enables the Network layer to determine a unique network address for a node. Routing protocols allow data to flow between networks and reach their proper destination. Examples of Network-layer protocols are Address Resolution Protocol (ARP), Datagram Delivery Protocol (DDP), Internet Control Message Protocol (ICMP), Interior Gateway Protocol (IGP), Internet Protocol (IP), Internetwork Packet Exchange (IPX), and Packet Layer Protocol (PLP). See also *ARP*, *DDP*, *ICMP*, *IGP*, *IP*, *IPX*, *OSI Reference Model*, *routing*.

network number—The portion of an IP address that denotes the network on which a host resides. The class of an IP address determines which portion of the address belongs to the network number and which portion belongs to the host number. See also *host number*, *IP address*.

Network-To-Network Interface—See *NNI*.

Network Virtual Terminal—See *NVT*.

NFS—Network File System. NFS is an Application-layer protocol, developed by Sun Microsystems, for sharing and transferring remote files on UNIX or other types of networks. See also *Application layer*.

NIC—See *InterNIC*.

NIC—Network Interface Card. A NIC enables a PC to connect to a network. The NIC uses drivers to communicate with the host's networking software, and interacts with the physical media that connects the host to other computers. A NIC is also called a *LAN adapter*, *network adapter*, or *network board*.

NIS—Network Information Service. Along with the Network File System (NFS), the NIS is a method of creating a distributed database system in order to centralize common configuration

files, such as the UNIX password file (`/etc/passwd`) and the hosts file (`/etc/hosts`). An NIS server manages copies of the database files, and NIS clients request information from them. NIS was developed by Sun Microsystems. See also *NFS*.

NNI—Network-To-Network Interface. A standard that defines the interface between two Frame Relay switches located in a private or public network. Both switches must reside in the same type of network. The User-To-Network Interface (UNI) standard defines the interface between a public switch and private one. Compare with *UNI*. See also *Frame Relay*.

node—See *host*.

node number—A value assigned to a host on a network. The node number can be hardcoded in the Network Interface Card (NIC), or assigned by means of jumper settings. It is unique amongst all the hosts on a local, physical network. The address for a host also contains the network address shared by all the hosts on the local network. See also *host*, *network address*, *NIC*.

Nonvolatile Random Access Memory—See *NVRAM*.

normal area—An Open Shortest Path First (OSPF) area that allows Type-5 Link-State Advertisements (LSAs) to be flooded throughout it. Area Border Routers (ABRs) advertise external routes as Type-5 LSAs. A normal area is the default for the Pipeline 220. If you change the default for one interface on the unit, you must change it for all interfaces, because the Pipeline 220 does not currently perform ABR functions. Compare with *NSSA*, *stub area*. See also *ABR*, *area*, *ASE Type-5*, *external route*, *LSA*, *OSPF*, *router*, *routing*.

Not So Stubby Area—See *NSSA*.

NSSA—Not So Stubby Area. An NSSA is an Open Shortest Path First (OSPF) area that does not receive or originate Type-5 Link-State Advertisements (LSAs), and that imports Autonomous System (AS) external routes in a limited fashion. OSPF version 2 defines a new Type-7 LSA for NSSAs. A Type-7 LSA differs from a Type-5 LSA in the following ways:

- An NSSA can originate Type-7 LSAs, and can advertise them throughout the NSSA.
- Type-7 LSAs are advertised only within a single NSSA. They are not flooded throughout the AS like Type-5 LSAs.

The Pipeline 220 can import ASE Type-7s only from static route definitions.

Compare with *normal area*, *stub area*. See also *area*, *AS*, *ASE Type-5*, *ASE Type-7*, *external route*, *LSA*, *OSPF*.

numbered interface—In interface-based IP routing, a unique address assigned to one side of a connection. When you use numbered interfaces, a local interface supports multiple IP addresses. One address is assigned in the interface's default profile, and one or more additional addresses are used for specific numbered-interface connections. Reasons for using numbered interfaces include troubleshooting nailed-up 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 220 to operate more as a multi-homed Internet host behaves. Compare with *system-based routing*, *unnumbered interface*. See also *interface-based routing*, *IP address*, *multi-homed host*, *point-to-point link*.

NVRAM—Nonvolatile Random Access Memory. NVRAM is a type of memory that maintains its data contents across resets and power cycles. It is useful for storing configuration information across sessions. Data is written and erased in blocks, rather than byte-by-byte.

The Pipeline 220 unit's system configuration is stored in the onboard NVRAM. Some error conditions may require that you clear the Pipeline 220 configuration and reboot. When you clear NVRAM, the system is re-initialized and comes up unconfigured, just as it was when you first installed it. You can then restore the configuration from a recent backup. NVRAM is also called *flash memory*. Compare with *DRAM*, *EEPROM*, *RAM*.

NVT—Network Virtual Terminal. An NVT is a bidirectional character device with a printer and a keyboard. The printer responds to incoming data, and the keyboard produces outgoing data sent over a Telnet connection. The code set is seven-bit ASCII in an eight-bit field. See also *NVT ASCII*, *Telnet session*.

NVT ASCII—The ASCII character set used with a Network Virtual Terminal (NVT). See also *ASCII*, *NVT*.

Octet—Eight data bits, also called a *byte*.

Open Shortest Path First—See *OSPF*.

Open Systems Interconnection Reference Model—See *OSI Reference Model*.

OSI Reference Model—Open Systems Interconnection Reference Model. The OSI Reference Model describes the layers of a network, details the functions of each layer, and explains how to connect communications devices on a LAN or WAN. Each layer provides services for the layer above it, and uses the services of the layer below it. The seven layers of the OSI model are as follows:

OSI Layer	Description
Application	Provides applications with access to the network. File transfer, email, and network management software are examples of Application-layer programs. Protocols such as Simple Network Management Protocol (SNMP), Telnet, Rlogin, File Transfer Protocol (FTP), and File Transfer, Access, and Management (FTAM) provide Application-layer services.
Presentation	Responsible for presenting information in a format understandable to users and their applications. Data conversion, special graphics, compression, and encryption are some of the functions implemented at the Presentation layer.
Session	Synchronizes the data in a network connection, maintains the link until the transmission is complete, handles security, and makes sure that the data arrives in the proper sequence. Gateway communications are implemented at the Session layer. Examples of Session-layer protocols are AppleTalk Data Stream Protocol (ADSP), NetBEUI (an extension of NetBIOS), NetBIOS, and Printer Access Protocol (PAP).
Transport	Provides data transfer at the proper speed, quality, and error rate, ensuring reliable delivery. Examples of Transport-layer protocols are Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and Sequenced Packet Exchange (SPX).
Network	Provides address resolution and routing protocols. Address resolution enables the Network layer to determine a unique network address for a node. Routing protocols allow data to flow between networks and reach their proper destination. Examples of Network-layer protocols are Address Resolution Protocol (ARP), Datagram Delivery Protocol (DDP), Internet Control Message Protocol (ICMP), Interior Gateway Protocol (IGP), Internetwork Packet Exchange (IPX), Internet Protocol (IP), and Packet Layer Protocol (PLP).
Data Link	Creates, sends, and receives data packets appropriate for the type of network in use. Data Link-layer protocols include High-Level Data Link Control (HDLC), Link Access Procedure, Balanced (LAPB), Link Access Procedure, D channel (LAPD), Point-To-Point Protocol (PPP), and Serial Line Internet Protocol (SLIP).
Physical	Defines the electrical properties of the physical medium, and converts the data into a series of 0s and 1s for digital transmission. Examples of Physical-layer specifications include RS-232, RS-422, RS-423, RS-449, IEEE 802.3, and IEEE 802.5.

OSPF—Open Shortest Path First. OSPF is the next generation Internet routing protocol. The *Open* in its name refers to the fact that OSPF was developed in the public domain as an open specification. The *Shortest Path First* portion refers to an algorithm developed by Dijkstra in 1978 for building a self-rooted shortest-path tree from which routing tables can be derived. As a link-state protocol, OSPF takes into account a variety of link conditions, such as the reliability or speed of the link, when determining the best path to a destination network. OSPF uses a link-state database of the network and propagates only changes to the database. See also *link-state database*, *route*, *router*, *routing*.

out-of-band management—A management method that uses a separate channel for diagnostic and administrative purposes (rather than a portion of each data channel).

output filter—A filter applied to an outgoing packet. See also *filter*, *packet filter*, *route filter*.

packet—A block of information containing a header, data, and trailer. Packets created at one level of the OSI Reference Model are inserted into lower-level packets. The format of a packet depends upon the protocol that creates it. A packet can be transmitted over a network or phone line. Compare with *frame*. See also *OSI Reference Model*, *packet field*.

packet field—A portion of a packet that contains a specific kind of information. For example, the data field in a packet contains the data being transmitted between applications. The header field can contain information identifying the packet type and any error-checking mechanisms. See also *packet*.

packet filter—A series of rules that instructs the Pipeline 220 on what to do when it encounters different types of packets. When you specify a packet filter, the Pipeline 220 monitors the data stream and takes a specified action when packet contents match the filter rules. Each filter specification either forwards or drops packets.

A match occurs at the first successful comparison between a filter and the packet being examined. When a comparison succeeds, the filtering process stops and the Pipeline 220 applies the forward action in the rule to the packet. If no comparisons succeed, the packet does not match the filter. However, the Pipeline 220 does not automatically forward packets. When no filter is in use, the Pipeline 220 forwards all packets, but once you apply a filter to an interface, the system *reverses* this default. For security purposes, the unit does not automatically forward non-matching packets. It requires a rule that explicitly allows those packets to pass.

You can apply a packet filter to incoming packets, outgoing packets, or both. In addition, you can specify that the Pipeline 220 forward or drop those packets that match the rules, or all packets *except* those that match the rules.

The Pipeline 220 supports two types of packet filters: generic and IP. You can apply a generic or IP filter as either a data filter or a call filter. The Pipeline 220 applies a data filter before a call filter.

Compare with *route filter*. See also *data filter*, *generic filter*, *IP filter*.

packet switching—A mode of data transfer in which packets are transmitted from a specific source to a specific destination using any available circuit. Packets may take different paths at the same time, and may not arrive in the order in which they were sent..

PAP—Password Authentication Protocol. PAP uses a two-way handshake method of establishing a caller's identity. Used only during the initial establishment of the data link, PAP is not a strong authentication method. Passwords travel across the line as plain text, so they are subject to eavesdroppers using software that monitors network information. Use PAP authentication only when the dial-in device does not support a stronger authentication method, such as Challenge Handshake Authentication Protocol (CHAP), or when the remote device requires a plain text password.

An extension of PAP adds the U.S. Data Encryption Standard (DES) cipher to data transmissions. The caller applies the encryption algorithm to a PPP packet and places the resulting cipher text in the information field of another PPP packet. The receiving end applies the inverse algorithm and interprets the resulting plain text as if it were a PPP packet that had arrived directly on the interface.

Compare with *CHAP*. See also *authentication*, *DES*, *password*, *PPP*.

password—A text string that a user must enter during the login process. Entering the proper password identifies the user as a person authorized to access network resources.

PBX—Private Branch Exchange. A PBX is an internal telephone network, such as those used in large offices, in which one incoming number directs calls to various extensions and from one office to another. See also *private circuit*.

PCMCIA—Personal Computer Memory Card International Association. PCMCIA is a standard that supports the devices on a credit-card-sized board. The 1990 PCMCIA version 1.0 specification supports Type I cards for RAM, ROM, or NVRAM. The 1991 PCMCIA version 2.01 specification supports Type II cards for network and fax/modem functionality, and Type III cards. A Type III card provides a miniature hard drive for wireless networks. See also *NVRAM*, *RAM*, *ROM*.

PCMCIA card—A card on the shelf controller that contains code for the slot cards, shelf-controller run-time code, and profiles. A PCMCIA card is also called a *flash card*. See also *PCMCIA*, *PCMCIA card code*, *PCMCIA slot*.

PCMCIA card code—Code written to make use of PCMCIA-card functionality. See also *PCMCIA*, *PCMCIA card*, *PCMCIA slot*.

PCMCIA slot—A slot on the Pipeline 220 shelf controller into which you can install a PCMCIA card. The Pipeline 220 contains two PCMCIA slots. See also *PCMCIA*, *PCMCIA card*, *PCMCIA card code*.

PDU—Protocol Data Unit. A PDU is a packet created at any one of the OSI layers. See also *OSI Reference Model*.

peripheral—A device attached to a network, server, or workstation. Peripherals include CD-ROM drives, fax machines, hard drives, modems, optical drives, printers, and tape drives.

permission level—A specification that governs the commands you can use at the Pipeline 220 Command-Line Interface (CLI). You set permission levels in a User profile. See also *User-To-Network Interface*.

Personal Computer Memory Card International Association—See *PCMCIA*.

per-user default route—The default route for IP packets coming from a particular user. The Pipeline 220 uses the per-user default under either of the following circumstances:

- The next-hop address in the Pipeline 220 unit's routing table is the default route for the system (destination 0.0.0.0).
- The normal routing logic fails to find a route and there is no system-wide default route. The direct route can take place by means of a profile or an Ethernet connection. If the Ascend unit does not have a direct route, it drops the packets on the connection. The default value is 0.0.0.0. If you accept this value, the Ascend unit routes packets as the routing table specifies, using the system-wide default route if it cannot find a more specific route. The per-user default route applies to all packets the Pipeline 220 receives for a given profile, regardless of the specific IP source address. Therefore, you can use this feature when the profile belongs to another router, and all hosts behind that router use the default gateway. The Pipeline 220 handles packets from other users or from the Ethernet network in the usual fashion. The global routing table is not altered. Therefore, when you diagnose routing

problems with a profile that implements a per-user default route, an error in a per-user gateway address is not apparent from inspection of the global routing table.
See also *default route, hop, IP address, IP route, IP routing table*.

Physical layer—The lowest layer in the OSI Reference Model. The Physical layer defines the electrical properties of the physical medium, and converts the data into a series of 0s and 1s for digital transmission. Examples of Physical-layer specifications include RS-232, RS-422, RS-423, RS-449, IEEE 802.3, and IEEE 802.5. See also *802.3, 802.5, OSI Reference Model, RS-232*.

Ping—A command that sends an Echo request in order to test whether a remote network device is accessible. If the remote device is properly connected, it receives the request and sends back an Echo reply. Certain version of the Ping command can also determine the amount of time necessary to receive the Echo reply, and the number of replies lost in transmission. See also *Echo*.

Point of Presence—See *POP*.

point-to-point link—A connection that does not make any use of intervening devices. A point-to-point link can connect two hosts on the same network, or two networks across the WAN.

Point-To-Point Protocol—See *PPP*.

POP—Point of Presence. A POP is the location of an Internet Service Provider's (ISP's) equipment. See also *ISP*.

port—A TCP/IP interface that defines the logical location in a computer where an application or process is running. When you define such a location, packets can reach an application from a remote system. There are certain well-known ports, such as port 21 used by FTP. Packet filters and firewalls make use of port addresses to restrict incoming and outgoing data and to secure an environment. The User Datagram Protocol (UDP) was developed to add the port address of an application or process to an IP packet, facilitating communication between applications over a network. See also *packet filter, firewall, IP, TCP/IP, UDP*.

POST—Power-On Self Test. A POST is a diagnostic test the Pipeline 220 performs when it first starts up or after it completes a system reset. During a POST, the Pipeline 220 checks system memory, configuration, installed cards, and T1 connections.

Power-On Self Test—See *POST*.

PPP—Point-To-Point Protocol. PPP provides a standard means of encapsulating data packets sent over a single-channel WAN link. It is the standard WAN encapsulation protocol for the interoperability of bridges and routers. PPP also allows direct dial-up access from a personal computer to a corporate LAN or Internet Service Provider (ISP). Using PPP ensures basic compatibility with non-Ascend devices. Both the dialing side and the answering side of the link must support PPP.

Typically, a dial-in device such as a modem or V.120 Terminal Adapter (TA) initiates a PPP session. The Pipeline 220 unit's terminal-server software handles the call. If the terminal server detects a PPP packet from the caller, it passes the call on to the router, which handles it as a regular PPP connection. The caller never sees the terminal-server interface.

However, if the user's dial-in software does not support PPP, the user can still initiate a PPP session from within the terminal-server interface. To do so, a user can log into the terminal server in terminal mode and use the PPP command. Or, you can include the PPP command in an expect-send script.

During establishment of a PPP data link, the dialing and answering units exchange Link Control Protocol (LCP) packets to establish communications and configure the link. When the link is established, PPP provides for an optional authentication step before exchanging Network Control Protocols (NCPs).

See also *bridge*, *ISP*, *LCP*, *NCP*, *router*, *terminal mode*, *terminal server*.

preference—A way for the Pipeline 220 to decide which route takes highest priority.

Routing Information Protocol (RIP) is a distance-vector protocol, which uses a hop count to select the shortest route to a destination network. Open Shortest Path First (OSPF) is a link-state protocol, which can take into account a variety of link conditions, such as the reliability or speed of the link, when determining the best path to a destination network. Because the metrics used by the two protocols are incompatible, the Pipeline 220 supports route preferences.

By default, static routes and RIP routes have the same preference, so they compete equally. Internet Control Message Protocol (ICMP) Redirects take precedence over both, and OSPF takes precedence over everything. If a dynamic route's preference is lower than that of the static route, the dynamic route can temporarily hide a static route to the same network.

However, dynamic routes age, and if no updates are received, they eventually expire. In that case, the hidden static route reappears in the routing table.

See also *dynamic route*, *hop count*, *ICMP*, *metric*, *OSPF*, *RIP*, *route*, *static route*.

Presentation layer—The second highest layer in the OSI Reference Model. The Presentation layer is responsible for presenting information in a format understandable to users and their applications. Data conversion, special graphics, compression, and encryption are some of the functions implemented at the Presentation layer. See also *OSI Reference Model*.

private circuit—See *nailed-up circuit*.

private network—A network particular to an organization, and not connected to a public data network such as the Internet. See also *VPN*.

profile—A collection of settings that enable you to configure various aspects of an Ascend product. For example, a Connection profile enables you to specify the name, password, and network resources for a dial-in caller. See also *Connection profile*.

profile index—See *index*.

Programmable Read-Only Memory—See *PROM*.

PROM—Programmable Read-Only Memory. PROM is a memory chip on which the system can write data only once. A PROM chip retains its contents across power cycles and system resets. See also *EEPROM*.

promiscuous mode—A bridging mode in which the Pipeline 220 unit's Ethernet controller accepts all packets and passes them up the protocol stack for a higher-level decision on whether to route, bridge, or reject them. Promiscuous mode is appropriate if you are using the Pipeline 220 as a bridge. See also *bridge*.

protocol—A set of rules governing message exchange over a network or internet. Examples of commonly used protocols are Transmission Control Protocol/Internet Protocol (TCP/IP), Point-To-Point Protocol (PPP), and Internetwork Packet Exchange (IPX). See also *internet*, *IPX*, *network*, *PPP*, *TCP/IP*.

Protocol Data Unit—See *PDU*.

proxy ARP—Proxy Address Resolution Protocol. Proxy ARP denotes a configuration in which one unit handles address resolution requests for another device. In an ARP request, a device asks a host to provide the host's physical address so that a connection can take place. ARP requests are broadcast only on the local network. If the Pipeline 220 is the default router on a network and is configured in proxy mode, packets destined for any of the hosts on the network go to the Pipeline 220. If a remote host must respond to an ARP request, the Pipeline 220 can respond on its behalf. See also *ARP*, *proxy mode*, *router*.

proxy mode—A mode in which a Connection profile assigns a local IP address to a remote host. Local hosts see the remote host as though it were on the local network. When calls are made to the remote host, the Pipeline 220 acts on its behalf, replying to requests and forwarding packets. See also *proxy ARP*.

public-key encryption—An encryption method that bases an encryption algorithm on the two halves of a long bit string. Each half of the bit sequence corresponds to a key. One key resides in a public-key library. Only a single party knows the other key. You can use either key to encrypt the data, but both keys are required to decrypt it. The sender can encrypt the data with the receiver's public key, and the receiver can decrypt it with the private key. Or, the sender can use private key to encrypt the message, and the receiver can use the public key to decrypt it. See also *encryption*.

PVC—Permanent Virtual Circuit. A PVC is a path maintained by two stations. The circuit is through the packet-switched network, but stays up all the time, regardless of whether or not data is on the line. Because the circuit is always up, there is no circuit setup time. Compare with *SVC*. See also *packet switching*.

RAM—Random Access Memory. RAM is computer memory that holds data temporarily. See also *DRAM*, *NVRAM*.

Random Access Memory—See *RAM*.

RARP—Reverse Address Resolution Protocol. RARP is a TCP/IP protocol that learns a workstation's hardware address and maps it to an IP address. See also *ARP*.

Raw 802.3—See *802.3*.

Raw TCP—Raw Transmission Control Protocol. Raw TCP is a method of supporting encapsulation performed by an application that runs on top of TCP. Raw TCP must be understood by both the login host and the caller. As soon as the connection is authenticated, the Pipeline 220 establishes a TCP connection to the host specified in the Connection profile. Raw TCP is also known as *TCP-Clear*.

Raw Transmission Control Protocol—See *Raw TCP*.

RBOC—Regional Bell Operating Company. An RBOC is one of seven companies created after the breakup of AT&T. The RBOCs are Ameritech, Bell Atlantic, Bell South, NYNEX, Pacific Telesis, Southwestern Bell, and U.S. West.

RDP—Reliable Data Protocol. RDP provides a reliable data transport service for packet-based applications. It is simple to implement, and works efficiently in environments that have long transmission delays and non-sequential delivery of message segments.

Read-Only Memory—See *ROM*.

redirect connection—A Frame Relay connection in which the Pipeline 220 ignores the destination IP address in a packet from a dial-in PPP client, and uses the Data Link Connection Indicator (DLCI) to route the packet instead. In effect, the Pipeline 220 does not route packets from the client in the usual sense. It simply passes them on to the Frame Relay network, and assumes that another device will route the packets on the basis of the destination IP address.

A Frame Relay redirect connection is not a full-duplex tunnel between the PPP dial-in device and the switch. The Pipeline 220 router handles the IP packets coming back from the Frame Relay switch, so the packets must contain the PPP caller's IP address for proper routing back across the WAN.

Compare with *circuit connection*, *gateway connection*. See also *DLCI*, *Frame Relay connection*, *IP address*, *PPP*.

redundancy—A method of safeguarding against line and equipment failure during a transmission. Each method for transmitting signals has inherent error rates, and all physical media is subject to damage. In the event of hardware failure, a redundant line or unit can take over at any time. You should always have a redundant (backup) module for multiplexers and other critical equipment.

reject interface—An interface that enables the router to handle packets whose IP address matches an unused IP address in a summarized address pool. The reject interface has an IP address of 127.0.0.2. When you specify this address as the router to the destination pool network, the Pipeline 220 rejects packets to an invalid host on that network, appending an ICMP Host Unreachable message. See also *POP*.

Reliable Data Protocol—See *RDP*.

remote device—A unit that resides across the WAN.

remote LAN Access—The process of allowing branch offices, telecommuters, and traveling computer users to access the corporate LAN backbone over digital or analog lines. The lines can be switched or nailed up. See also *digital line*, *nailed-up line*.

remote network—A network to which the Pipeline 220 connects over the WAN.

Remote Procedure Call—See *RPC*.

remote user—A user at a device not connected directly to the Ascend unit and not residing on the local Ethernet.

replay attack—A strategy for gaining illegal access to a system. During a replay attack, an unauthorized user records valid authentication information exchanged between systems, and

then replays it later to gain entry. Token-card authentication protects your system against replay attacks. Because the token is a one-time-only password, replay is impossible.

Request For Comments—See *RFC*.

Reverse Address Resolution Protocol—See *RARP*.

RFC—Request for Comments. RFC denotes the document series, begun in 1969, that describes the Internet suite of protocols and related experiments. Not all RFCs describe Internet standards, but all Internet standards are written up as RFCs. The RFC series of documents is unusual in that the proposed protocols are distributed by the Internet research and development community, acting on its own behalf. The protocols do not go through the formal review and standardization process promoted by organizations such as ANSI. A complete list of RFCs resides at <http://www.internic.net/rfc/>.

RIP—Routing Information Protocol. RIP is a distance-vector protocol found in both the NetWare and TCP/IP protocol suites. The protocol keeps a database of routing information that it gathers from periodic broadcasts by each router on a network. See also *distance-vector metric*, *router*, *routing*, *TCP/IP*.

ROM—Read-Only Memory. ROM is computer memory whose contents can be read and executed, but not modified. See also *EEPROM*, *PROM*.

route—The path that data takes from its source network to its destination network. See also *IP route*, *IPX route*.

router—A device that determines a path from a host on one network to a host on another. The networks may be in close proximity, or may be separated by long distances. A router has access to the three lowest OSI layers, and generally operates at the Network layer. To route a packet, a router uses the logical address specified as the packet's destination field, and determines the next router (if any) to which the packet must travel to reach its destination. All routers share information about the current topology and state of the network, maintaining routing tables that reflect the latest information. See also *IP router*.

router mode—An Ascend Tunnel Management Protocol (ATMP) configuration in which the home agent's routing module forwards packets it receives from the foreign agent onto the local network. The network can be the home network, or it can support another router that can connect to the home network. In either case, packet delivery relies on a routing mechanism, such as a static or dynamic route, and not on a WAN connection. Compare with *gateway mode*. See also *ATMP*, *dynamic route*, *foreign agent*, *home agent*, *home network*, *static route*.

route filter—A type of filter containing rules for the action to take on routes in Routing Information Protocol (RIP) update packets. When you apply a route filter to an IP interface, the Pipeline 220 monitors RIP packets on the interface and takes one of the following actions when a route matches the filter rules:

- No action (the default).
 - Accept the route by allowing it to affect the routing table.
 - Deny the route by not allowing it to affect the routing table.
 - Add the value set in the Add-Metric parameter to the route metric and accept the route.
- The filter can apply to incoming packets, outgoing packets, or both. When you apply a route filter to an interface, the Pipeline 220 applies all defined input and output filters to RIP update

packets until it finds a match. If it does not find a match for a route, the default action is to deny the route. Compare with *packet filter*. See also *IP route*, *IP routing table*, *RIP*.

routing—A method of determining how to forward a data packet to the proper destination. See also *IP routing*, *IPX routing*, *OSPF*, *RIP*, *route*, *router*.

Routing Information Protocol—See *RIP*.

routing table—See *IP routing table*.

RPC—Remote Procedure Call. An RPC is a method in which a program on one device can transparently use a procedure on another device. RPCs are often used in client-server architectures.

RS-232—An EIA standard that specifies various electrical and mechanical characteristics for interfaces between Data Terminal Equipment (DTE) and Data Circuit-Terminating Equipment (DCE) devices. The standard applies to both synchronous and asynchronous binary data transmission at rates below 64 Kbps. RS-232 is also known as *EIA/TIA-232*. See also *DCE*, *DTE*, *synchronous transmission*.

SAM—Secure Access Manager. SAM gives you a high degree of centralized control over the security functions of an entire network. Through this Windows-based application, you can configure Secure Access Firewall(s) offline, and download the configuration to remote locations. See also *Secure Access Firewall*.

SAP—Service Access Point. A SAP is a defined location through which a procedure at one OSI layer can provide services to the next layer above it. Each SAP has a unique address in hexadecimal format. See also *DSAP*, *OSI Reference Model*, *SSAP*.

SAP—Service Advertising Protocol. SAP is a NetWare protocol that operates at the Transport layer and enables servers to inform other devices about the services they have available. Each server advertises its services using a SAP packet. Each router on the network retrieves the SAP packets and builds a database of all the servers it knows about. The router then broadcasts this information to other routers, either at a set interval or whenever the database changes. See also *IPX router*, *IPX server*.

SAP filter—Service Advertising Protocol filter. A filter that determines which SAP advertisements the Pipeline 220 forwards or drops. The router examines incoming and outgoing SAP packets to see whether certain fields in the packet match the filter. A SAP filter enables you to control the size of resident SAP tables and reduce bandwidth usage. You can also use a SAP filter to restrict a user's view of services on the network.

SDRP—Source Demand Routing Protocol. SDRP supports source-initiated selection of interdomain routes, working along with the intermediate node selection provided by Border Gateway Protocol (BGP) and Inter-Domain Routing Protocol (IDRP). See also *BGP*, *IDRP*.

Secure Access Firewall—An Ascend software option that stops intruders from breaking and entering into your network. A firewall is similar to a filter, but is more complex, dynamically changing in response to the characteristics of the packets that pass through it. The firewall affects which packets are allowed to reach the network, and which packets can leave the network for another interface. Typically, you can design a firewall to flag a packet with

specific bit patterns, and put rules into action that cause other rules to be created. For a firewall to take effect, you must apply it to a LAN or WAN interface. See also *SAM*.

Secure Access Manager—See *SAM*.

serial communication—Communication through the serial port of a device. For Windows 3.1, the maximum speed of the serial port is 19,200. For Windows 95, the serial port limit is 921,600. These limitations are subject to change with the development of a faster serial bus. See also *serial port*, *serial transmission*.

serial connection—A link between the serial ports of two devices. See also *serial communication*, *serial port*, *serial transmission*.

serial host—A device (such as a videoconferencing codec) that is connected to a serial WAN port communicating over a point-to-point link. To a serial host, the Pipeline 220 appears to be a cable or Data Circuit-Terminating Equipment (DCE). See also *DCE*, *point-to-point link*, *serial WAN port*.

serial port—A port that transmits and receives asynchronous or synchronous serial data. See also *serial transmission*, *synchronous transmission*.

serial transmission—A form of data transmission in which only one line carries all eight bits of a byte. In serial transmission, one bit follows another (as opposed to parallel transmission, in which the bits travel simultaneously, each on a different wire). Serial transmission can be either synchronous or asynchronous. Synchronous communication requires additional lines for transmitting handshake or timing signals. In asynchronous communication, the data itself contains synchronization information, so neither handshake nor clock signals are necessary. See also *synchronous transmission*.

serial WAN port—A port that provides a V.35/RS-449/X.21 WAN interface, typically used to connect the Pipeline 220 to a Frame Relay switch. The clock speed received from the link determines the serial WAN data rate. The maximum acceptable clock is 8 Mbps. The clock speed at the serial WAN port has no effect on the bandwidth of other WAN interfaces in the Pipeline 220. See also *serial transmission*.

server—A device or program that provides services to hosts on a network.

Service Access Point—See *SAP*.

Service Advertising Protocol—See *SAP*.

session—The state a connection reaches when two parties can communicate with each other.

session ID—A unique ID that denotes a particular Pipeline 220 session. The Pipeline 220 can pass a session ID to SNMP, or other external entities. See also *session*, *session ID base*.

session ID base—The base number for calculating a session ID. If the value of the session ID base is nonzero, the Pipeline 220 uses it as the initial base for calculating session IDs after a system reset. The system increments the ID for each subsequent session by 1. If the session ID base is zero, the Pipeline 220 sets the initial base for session IDs to the absolute clock. For example, if the clock is 0x11cf4959, the subsequent session IDs uses 0x11cf4959 as a base.

However, if the clock changes and the system reboots or clears NVRAM, session IDs may be duplicated. See also *session*, *session ID*.

Session layer—The third highest layer in the OSI Reference Model. The Session layer synchronizes the data in a network connection, maintains the link until the transmission is complete, handles security, and makes sure that the data arrives in the proper sequence. Gateway communications are implemented at the Session layer. Examples of Session-layer protocols are AppleTalk Data Stream Protocol (ADSP), NetBEUI (an extension of NetBIOS), NetBIOS, and Printer Access Protocol (PAP). See also *OSI Reference Model*.

Shielded Twisted Pair cable—See *STP cable*.

Signaling System 7—A protocol architecture that specifies a series of Signaling Points (SPs) and Signaling Transfer Points (STPs) connected on a network. The SPs are hosts from which signaling messages originate and terminate. The STPs are packet switches that perform message routing between adjacent SPs or STPs. The Network Services Part (NSP) of the Signaling System 7 provides reliable message transfer, and corresponds to the lower three layers of the OSI model. The NSP consists of a Message Transfer Part (MTP) and a Signalling Connection Control Part (SCCP). See also *OSI Reference Model*.

Simple Mail Transfer Protocol—See *SMTP*.

Simple Network Management Protocol—See *SNMP*.

Simple Network Time Protocol—See *SNTP*.

slot compression—Compression in which the slot ID does not appear in any VJ-compressed packet but the first in the data stream. When you turn on VJ compression, the Pipeline 220 removes the TCP/IP header, and associates a TCP/IP packet with a connection by giving it a slot ID. The first packet coming into a connection must have a slot ID, but succeeding packets need not have one. By default, the Pipeline 220 uses slot compression: if the packet does not have a slot ID, the Pipeline 220 associates it with the last-used slot ID. See also *VJ compression*.

SMDS—Switched Multimegabit Data Service. SMDS is a packet-based service that enables the creation of high-speed data networks with rates of up to 45 Mbps.

SMTP—Simple Mail Transfer Protocol. In the TCP/IP protocol suite, SMTP is an Application-layer protocol that uses the TCP Transport-layer protocol to send and receive electronic mail. See also *TCP/IP*.

SNAP—SubNetwork Access Protocol. SNAP is a protocol specification for the format of the Media Access Control (MAC) header of an IPX frame. SNAP includes the IEEE 802.3 protocol format plus additional information in the MAC header. Compare with *802.2*, *802.3*, *Ethernet II*. See also *IPX frame*, *MAC*.

SNMP—Simple Network Management Protocol. SNMP is a standard way for computers to share networking information. In SNMP, two types of communicating devices exist: agents and managers. An agent provides networking information to a manager application running on another computer. The agent can be polled by the manager, and can also use a message called a traps-PDU to send unsolicited information to the manager when an unusual event occurs. The Pipeline 220 is an example of

an SNMP agent. The agents and managers share a database of information, called the Management Information Base (MIB).

The Pipeline 220 supports SNMP MIB II, T1 MIB, and Ascend Enterprise MIBs. A manager that uses the Ascend Enterprise MIB can query the Pipeline 220, set parameters, sound alarms when certain conditions appear, and perform other administrative tasks. You can therefore manage the Pipeline 220 from a central SNMP manager, such as SunNet Manage or HP Open View.

SNMP security uses the community name that the manager sends (with each polling request) and that the agent sends (with each traps-PDU). Ascend supports two community names: one with read-only access, and the other with read/write access to the MIB.

SNTP—Simple Network Time Protocol. SNTP enables a server to retrieve the correct time from an official source and distribute the information to other servers and networks. The protocol also enables a group of servers to synchronize their clocks with reference to a primary time server. See also *UTC*.

socket—A TCP/IP interface that facilitates a two-way link between systems, enabling applications to run over a connectionless network. A socket is defined by two addresses: the IP address of the host computer, and the port address of the application or process running on the host. See also *IP address, port, TCP/IP*.

socket number—A unique value assigned to a socket in a network. See also *socket*.

soft IP interface address—An IP address that is not associated with a physical interface. A soft interface is just like any other interface on the Pipeline 220, except that it is never down. In general, the Pipeline 220 uses the soft IP address for incoming packets. You can also specify the soft interface address as the system IP address, in which case it becomes the source address for the traffic generated by the Pipeline 220. Routing protocols advertise the address as a host route with a mask of 32, using the loopback interface.

Other hosts on the network must be able to reach the address you assign as the soft interface address, so you must either enable routing protocols (RIP, OSPF) or configure static routes in routers one hop away from the Pipeline 220. To verify that other hosts in your network have a route to the soft address, use Ping and Traceroute from the other hosts to the Pipeline 220.

Compare with *system-based routing*. See also *IP address, OSPF, RIP, static IP route*.

software compression—See *compression*.

Source Demand Routing Protocol—See *SDRP*.

Source Service Access Point—See *SSAP*.

SSAP—Source Service Access Point. An SSAP is the Service Access Point (SAP) address at which at a Network-layer procedure requests services from the Logical Link Control (LLC) layer. See also *DSAP, SAP*.

Stac Lempel-Ziv standard compression—See *Stac LZS compression*.

Stac compression—On the Pipeline 220, a compression option that specifies an Ascend-modified version of draft 0 of the CCP (Compression Control Protocol). The Stac option is an Ascend variant of the Stac LZS compression method. It was implemented before Stac LZS was standardized. Compare with *Stac LZS compression*.

Stac-9 compression—On the Pipeline 220, a compression option that indicates the method specified by draft 9 of the Stac LZS compression protocol. Compare with *Stac compression*. See also *Stac LZS compression*.

Stac LZS compression—Stac Lempel-Ziv standard compression. Developed by Stac Incorporated, Stac LZS compression can triple data rates. Compare with *Stac compression*. See also *Stac-9 compression*.

start bit—In asynchronous transmission, a bit that indicates the beginning of a new character. It is always 0 (zero).

static IP route—A path that specifies a destination IP network and the gateway (next-hop router) to get to that network. Each Connection profile that specifies an explicit IP address defines a static route to a remote or local IP network. Compare with *dynamic route*, *multipath route*. See also *Connection profile*, *IP address*, *IP network*.

static IPX route—A route (configured in an IPX-Interface profile) that contains all the information necessary to reach one IPX server on a remote network. The Pipeline 220 adds the static routes upon initialization. When the Pipeline 220 receives an outgoing packet for a server, it finds the corresponding profile and dials the connection. You must manually update static routes whenever the administrator at the remote end removes the specified server or updates its address. You do not need to create IPX routes to servers that reside on the local Ethernet network. See also *IPX server*.

static password—A password specified in a Connection profile. The user must enter the password to gain access to the Pipeline 220. See also *Connection profile*.

static route—See *static IP route*, *static IPX route*.

station—See *host*.

status window—A window in the Pipeline 220VT100 interface that displays system status information. The default status window contains three window areas—a large portion on the left, and a portion on the right consisting of top and bottom windows. The left side displays WAN connection and session status. The right-hand windows can displays general status information, data on Ethernet activity, the contents of the log buffer, and line statistics.

STP cable—Shielded Twisted Pair cable. STP cable consists of two wires twisted two or more times per inch in order to help cancel out noise. The entire cable has a protective covering. STP cable is typically used in ARCnet and Token Ring networks.

straight-through cable—A cable with wires that have terminating ends with the same wire assignments. Compare with *crossover cable*.

stub area—An Open Shortest Path First (OSPF) area in which all external routes are summarized by a default route. To reduce the cost of routing, OSPF supports stub areas. A stub area allows no Type-5 LSAs to be propagated in the area. Instead, it depends on default routing to external destinations. Compare with *normal area*, *NSSA*. See also *area*, *Open Shortest Path First*.

subnet—See *IP subnet*.

subnet mask—An IP feature in which a group of bits identifies a subnet. To specify a subnet mask, the Pipeline 220 appends to the IP address a modifier that specifies the total number of network bits in the address. For example, in the address 198.5.248.40/29, the /29 specification indicates that 29 bits of the address specify the network. The three remaining bits specify unique hosts. With three bits used to specify hosts on a 29-bit subnet, eight different bit-combinations are possible:

000—Reserved for the network (base address)

001

010

100

110

101

011

111—Reserved for the broadcast address of the subnet

The standard and Ascend subnet formats for a class C network number are as follows:

Standard subnet mask	Number of host addresses	Ascend notation
255.255.255.0	254 hosts + 1 broadcast, 1 network base	/24
255.255.255.128	126 hosts + 1 broadcast, 1 network base	/25
255.255.255.192	62 hosts + 1 broadcast, 1 network base	/26
255.255.255.224	30 hosts + 1 broadcast, 1 network base	/27
255.255.255.240	14 hosts + 1 broadcast, 1 network base	/28
255.255.255.248	6 hosts + 1 broadcast, 1 network base	/29
255.255.255.252	2 hosts + 1 broadcast, 1 network base	/30
255.255.255.254	Invalid subnet mask (no hosts)	/31
255.255.255.255	1 host (a host route)	/32

The broadcast address of any subnet has the host portion of the IP address set to all ones. The network address (or base address) represents the network itself, because the host portion of the IP address is all zeros. For example, suppose the Pipeline 220 configuration assigns the following address to a remote 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

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.

See also *host number*, *host route*, *IP*, *IP address*, *IP subnet*, *network number*.

SubNetwork Access Protocol—See *SNAP*.

SVC—Switched Virtual Circuit. An SVC is a path over a packet-switched network. It appears to be a dedicated circuit, but the connection stays up only as long as needed. Compare with *PVC*.

SWIPE—IP with Encryption. SWIPE is a Network-layer security protocol that works by adding a cryptographic authenticator to each packet, and encrypting the data.

switch—A device that connects the calling party to the answering party.

Switched Multimegabit Data Service—See *SMDS*.

symbolic name—A name that denotes an IP address. A symbolic name consists of a user name and a domain name in the format *username@domain_name*. The user name corresponds to the host number in the IP address. The domain name corresponds to the network number in the IP address. A symbolic name might be *steve@abc.com* or *joanne@xyz.edu*. See also *IP address*.

synchronization—A method of ensuring that the receiving end can recognize characters in the order in which the transmitting end sent them, and can know where one character ends and the next begins. Without synchronization, the receiving end would perceive data simply as a series of binary digits with no relation to one another.

synchronous transmission—A transmission mode in which the data moves in large blocks, called messages or frames. A synchronous WAN link uses High-level Data Link Control (HDLC) encoding and connects to a router for a network-to-network link. The Pipeline 220 routes a synchronous transmission as a digital call to an HDLC channel, and then to the router software. Each synchronous call uses Point-To-Point Protocol (PPP), Multilink Protocol (MP), Multilink Protocol Plus (MP+), or Frame Relay encapsulation.

In a synchronous transmission, both the sending device and the receiving device must maintain synchronization in order to determine where one block of data ends and the next begins. Each side can transmit a separate synchronizing signal, called a clock. Or, each frame or message can contain synchronization information.

In the latter method, each block of data starts with one or more control characters, usually eight bytes long, called a SYNC. The receiver interprets the SYNC as a signal that it can start accepting data. Synchronous transmission can be up to 20 percent faster than asynchronous transmission.

See also *Frame Relay*, *HDLC*, *PPP*, *synchronization*.

Syslog host—The station to which the Pipeline 220 sends system logs.

system-based routing—A form of IP routing in which the entire unit has a single IP address. For systems that have a single backbone connection, system-based routing is the simplest way to configure the Pipeline 220. Compare with *interface-based routing*.

system status window—See *status window*.

T1 channel—One of 24 channels on a T1 line. See also *fractional T1 line*, *T1 line*, *T1 PRI line*, *unchannelized service*.

T1 line—A line that supports 24 64-Kbps channels, each of which can transmit and receive data or digitized voice. The line uses framing and signaling to achieve synchronous and reliable transmission. The most common configurations for T1 lines are ISDN Primary Rate Interface (T1 PRI) and unchannelized T1, including fractional T1. See also *fractional T1 line*, *T1 channel*, *T1 PRI line*, *unchannelized service*.

T1 PRI line—T1 Primary Rate Interface line. A T1 PRI line has a total bandwidth of 1.544 Mbps. It uses 23 B channels for user data, and one 64-Kbps D channel for ISDN D-channel signaling. The B channels can be all switched, all nailed up, or a combination of switched and nailed up. The T1 PRI line is a standard in North America, Japan, and Korea. You can connect this type of line to standard voice, Switched-56, Switched-64, Switched-384, Switched-1536, and MultiRate data services. Using a feature called PRI-to-TI conversion, the Pipeline 220 can share the bandwidth of a T1 PRI line with a PBX. Compare with *E1 PRI line, unchannelized service*. See also *DDP, nailed-up channel, nailed-up channel, PBX, private circuit, T1 channel, T1 line*.

T1 Primary Rate Interface line—See *T1 PRI line*.

tag—An Open Shortest Path First (OSPF) method of flagging a route as external—that is, as having been imported into the OSPF database from outside the router's Autonomous System (AS). See also *AS, external route, OSPF*.

tariff—A document filed by a regulated telephone company with a state public utility commission or the Federal Communications Commission. A tariff details services, equipment, and pricing publicly offered by the telephone company.

TCP—Transmission Control Protocol. TCP operates at the Transport layer, providing connected-oriented services. It uses IP to deliver packets. See also *IP*.

TCP-Clear—See *Raw TCP*.

TCP/IP—Transmission Control Protocol/Internet Protocol. TCP/IP is a family of protocols that defines the format of data packets sent across a network, and is the communications standard for data transmission between different platforms. TCP/IP defines the following family of protocols and services:

TCP/IP Protocol name	Description of service
Transmission Control Protocol (TCP)	Transport protocols that control data transmission between computers
User Datagram Protocol (UDP)	
Internet Protocol (IP)	Routing protocols that control addressing and packet assembly, and determine the best route for a packet to take to arrive at its destination
Internet Control Message Protocol (ICMP)	
Routing Information Protocol (RIP)	
Open Shortest Path First (OSPF)	
Border Gateway Protocol version 4 (BGP)	Gateway protocols that enable networks to share routing and status information
Gateway-To-Gateway Protocol (GGP)	
Interior Gateway Protocol (IGP)	
Domain Name System (DNS)	Network-address services and protocols that handle the way each computer on a network is identified
Address Resolution Protocol (ARP)	
Reverse Address Resolution Protocol (RARP)	
Boot Protocol (BOOTP)	User services that provide applications a computer can use
File Transfer Protocol (FTP)	
Telnet	
Network File System (NFS)	File-transfer, mail, and management services
Network Information Service (NIS)	
Remote Procedure Call (RPC)	
Simple Mail Transfer Protocol (SMTP)	
Simple Network Management Protocol (SNMP)	

See also *Address Resolution Protocol*, *BGP*, *BOOTP*, *DNS*, *EGP*, *FTP*, *GGP*, *ICMP*, *IGP*, *IP*, *NFS*, *NIS*, *OSPF*, *RARP*, *RIP*, *RPC*, *SMTP*, *SNMP*, *TCP*, *Telnet*, *UDP*.

TCP/IP header compression—See *VJ compression*.

TDM—Time Division Multiplexing. TDM is a scheme that uses time-slot assignment, enabling information from multiple channels to use bandwidth on a single line.

Telecommunications Industry Association—See *TIA*.

telecommuter—A work-at-home computer user who connects to the corporate LAN backbone by means of remote-access technology. For example, a telecommuter can establish a link with the LAN by means of a modem connected to an analog line, an ISDN Terminal Adapter (TA) or router connected to an ISDN line, or a Channel Service Unit/Data Service Unit (CSU/DSU) connected to a Switched-56 line. See also *CSU*, *DSU*.

Telnet—A protocol that links two computers in order to provide a terminal connection to the remote machine. Instead of dialing into the computer, you connect to it over the Internet using Telnet. When you issue a Telnet session, you connect to the Telnet host and log in. The connection enables you to work with the remote machine as though you were a terminal connected to it. If your Pipeline 220 has an Ethernet card installed, you can remotely manage it by establishing a Telnet session to the Pipeline 220 from any Telnet workstation on the network.

Telnet session—A terminal connection to a remote machine by means of the Telnet protocol. After you set up a basic IP configuration for the Pipeline 220, users can Telnet into the Pipeline 220 command line. Each user can initiate a Telnet session to the Pipeline 220 from a local workstation or from a WAN connection. In both cases, the Pipeline 220 authenticates the session by means of a User profile, which defines a permission level for the user logging in. In addition to the password required by a User profile, you can specify that Telnet requires its own password authentication, which occurs prior to any User profile authentication. See also *User-To-Network Interface*.

terminal—A computer that does not have its own processor and that must connect to a terminal server in asynchronous mode to use its Central Processing Unit (CPU). VT100, ANSI, and TTY are all types of terminals.

terminal mode—A terminal-server access mode in which the Pipeline 220 negotiates a user-to-host session. Instead of providing only the login name and password required to authenticate a Connection profile, you can set up an expect-send script that also includes the terminal-server prompt and a command, such as PPP, SLIP, TCP, Telnet, or Rlogin. In this way, the session with a host comes about as part of the login process, so the user never actually sees the terminal-server command-line prompt. Alternatively, you can provide access to the command line and restrict the commands you make accessible to the user. See also *PPP*, *TCP*, *Telnet*, *terminal server*.

terminal server—A terminal server is a computing device to which a terminal can connect over a LAN or WAN link. A terminal communicates with the terminal server over an asynchronous serial port (typically an RS-232 port) through a modem. A terminal converts the data it receives from the terminal server into a display and does no further processing of the data. A terminal also converts the operator's keystrokes into data for transmission to the terminal server.

The Pipeline 220 terminal-server software receives asynchronous calls after they have been processed by a digital modem. Typically, a modem or V.120 Terminal Adapter (TA) dials these calls. V.120 and TCP calls are enabled by default. If the caller does not send Point-To-Point Protocol (PPP) packets immediately, the terminal server starts a login sequence.

Each user must have a Connection profile that specifies a name and password to use in the terminal-server login sequence. In addition, a global Terminal-Server profile defines how terminal-server calls are authenticated, and determines the destination of the call after authentication is complete. When it receives a name and password from the caller, the terminal server authenticates them by means of a Connection profile or external authentication server, and then performs one of the following actions:

- Displays the terminal-server command-line prompt
- Displays a menu of hosts the user can log into
- Immediately logs the user into a designated host
- Initiates a PPP or SLIP session with the user

To protect the command-line from unauthorized access, you can also choose to assign the terminal server its own password.

If it receives an asynchronous PPP call, the terminal server does not begin a login sequence. Instead, it responds with a PPP packet, and Link Control Protocol (LCP) negotiation begins, including negotiation for Password Authentication Protocol (PAP) or Challenge Handshake Authentication Protocol (CHAP) authentication. The terminal server directs the call to the router software, and the connection proceeds as for a regular synchronous PPP session. The user bypasses the terminal server interface altogether.

In most cases, the terminal server is a stepping stone toward access to one or more network hosts. To enable host access, you can configure the terminal server in terminal mode, immediate mode, or menu mode.

See also *CHAP*, *Connection profile*, *digital signal*, *menu mode*, *PAP*, *PPP*, *terminal mode*.

terminal-server connection—A connection between a terminal and a terminal server over a LAN or WAN link. See also *terminal server*.

terminal-server session—An end-to-end connection between a terminal and a terminal server. Usually, the terminal-server session begins when the call goes online and ends when the call disconnects. The Pipeline 220 supports all the common capabilities of standard terminal servers, including Telnet, Domain Name System (DNS), login and password control, Call Detail Reporting (CDR), and authentication services. See also *terminal server*.

Thick Ethernet—A type of .4" diameter coaxial cable for Ethernet networks. Also known as *thicknet*.

thicknet—See *Thick Ethernet*.

Thin Ethernet—A type of .2" diameter coaxial cable for Ethernet networks. Also known as *thinnet*.

thinnet—See *Thin Ethernet*.

third-party routing—A feature that enables the Pipeline 220 to advertise Open Shortest Path First (OSPF) routes to external destinations on behalf of another gateway, commonly known as advertising a forwarding address. When third-party routing is enabled, the Pipeline 220 advertises the IP address of another gateway. If third-party routing is disabled, the Pipeline 220 advertises itself as the forwarding address to an external destination.

Depending on the exact topology of the network, other routers might be able to route packets directly to the forwarding address without involving the advertising Pipeline 220, increasing the total network throughput. In this scenario, all OSPF routers must know how to route to the forwarding address.

See also *OSPF*.

TIA—Telecommunications Industry Association. The TIA is a group that determines standards for the electrical level of data transmission.

tick—An IBM unit of measurement that corresponds to one-eighteenth of a second.

Time Division Multiplexing—See *TDM*.

timeout—An event in which a device or user exceeded a configured time limit for responding to a device or process.

Transmission Control Protocol—See *TCP*.

Transmission Control Protocol/Internet Protocol—See *TCP/IP*.

transparent bridge—A bridge that notes a packet's source address and creates a bridging table that associates a host's Media Access Control (MAC) address with a particular Ethernet

interface. The Pipeline 220 is an example of a transparent bridge (also called a *learning bridge*). See also *bridge*, *bridging*.

Transport layer—The middle layer of the OSI Reference Model. The Transport layer provides data transfer at the proper speed, quality, and error rate, ensuring reliable delivery. Examples of Transport-layer protocols are Transmission Control Protocol (TCP) and User Datagram Protocol (UDP). See also *OSI Reference Model*, *TCP*, *UDP*.

traps-PDU—A message that a Simple Network Management Protocol (SNMP) agent sends to a manager application to inform the manager of network events. See also *agent*, *community name*, *manager*, *MIB*, *SNMP*.

tunneling—A way of overcoming protocol restrictions on a network by encapsulating packets that use one protocol inside packets that use a protocol supported by the network.

twisted-pair cable—A cable consisting of four or more copper wires twisted together in pairs. Telephone wiring is an example of twisted-pair cable. Twisted-pair cable can be shielded or unshielded. See also *STP cable*, *UTP cable*.

twisted-pair Ethernet—See *10Base-T*.

two-wire subscriber loop—A two-wire WAN link connecting the Customer Premises Equipment (CPE) to the carrier's switch. See also *CPE*.

UDP—User Datagram Protocol. UDP is a Transport-layer protocol that provides connectionless service without packet acknowledgment. See also *Transport layer*, *UDP port*.

UDP port—A 16-bit number that allows multiple processes to use User Datagram Protocol (UDP) services on the same host. A UDP address is the combination of a 32-bit IP address and the 16-bit port number. Examples of well-known UDP ports are 7 (for Echo packets), 161 (for SNMP packets), and 514 (for Syslog packets). See also *UDP*.

unchannelized E1—See *unchannelized service*.

unchannelized service—A service that uses the entire bandwidth of a T1 PRI line (1.544 Mbps) or an E1 PRI line (2.048 Mbps). You can use an unchannelized line for a nailed-up connection, such as the link to a Frame Relay network. The Pipeline 220 treats the line as though it were a single connection at a fixed speed, without individual channels. See also *E1 PRI line*, *T1 PRI line*.

unchannelized T1—See *unchannelized service*.

UNI—User-To-Network Interface. UNI is the interface between an end user and a network endpoint (a router or a switch) on the Frame Relay network. See also *UNI-DCE interface*, *UNI-DTE interface*.

unicast network—A network in which a router sends packets to one user at a time. Compare with *broadcast network*, *multicast network*.

UNI-DCE interface—User-To-Network Interface—Data-Circuit-Terminating-Equipment Interface. In a UNI-DCE configuration, the Pipeline 220 operates as a Frame Relay router

communicating with a Data Terminal Equipment (DTE) device. To the DTE device, the Pipeline 220 appears as a Frame Relay network endpoint.

When you set up a Pipeline 220 in a UNI-DCE configuration, it can perform DCE link management functions. The Pipeline 220 expects to get regular requests for the status of the link. If the Pipeline 220 does not receive the requests within the expected interval, it considers the link inactive. The Pipeline 220 responds to the requests with the status of the link identified by the Data Link Connection Indicator (DLCI). Compare with *UNI-DTE interface*. See also *DCE, DLCI, DTE, Frame Relay, UNI*.

UNI-DTE interface—User-To-Network Interface—Data-Terminal-Equipment Interface. In a UNI-DTE connection, the Pipeline 220 is a Data Terminal Equipment (DTE) device communicating with a Frame Relay switch. The Pipeline 220 acts as a Frame Relay feeder and can perform DTE functions for link management. When it performs DTE link management, the Pipeline 220 regularly requests updates on the status of the link. If the Frame Relay unit at the other end of the link does not respond to the requests, or if the response indicates a Data Link Connection Indicator (DLCI) failure, the Pipeline 220 considers the link inactive. In addition, the Pipeline 220 can query the device at the other end of the link about the status of the DLCIs in the connection. If any of the DLCIs becomes unusable and the DLCI's local Connection profile specifies a backup connection, the Pipeline 220 dials the connection. Compare with *UNI-DCE interface*. See also *DCE, DLCI, DTE, Frame Relay switch, UNI*.

Universal Time Configuration—See *UTC*.

UNIX—A 32-bit operating system that allows multiple users to share resources and perform multiple tasks at the same time. UNIX was developed at Bell Laboratories in 1969. Its development has occurred along two lines: the AT&T System versions and the UC Berkeley Distribution (BSD) releases. The two strains were combined by the UNIX Systems Group into System V Release 4.2 (SVR 4.2).

unnumbered interface—A link that uses system-based routing, in which the entire Pipeline 220 system has a single IP address. If all interfaces are unnumbered, the Pipeline 220 operates as a purely system-based router. Compare with *interface-based routing, numbered interface*. See also *IP routing, system-based routing*.

Unshielded Twisted Pair cable—See *UTP cable*.

User Datagram Protocol—See *UDP*.

user name—The name a user must enter to access the services of the Pipeline 220. See also *password*.

User-To-Network Interface—See *UNI*.

UTC—Universal Time Configuration. Formerly known as Greenwich Mean Time (GMT), UTC is the time at the Greenwich observatory, used as a reference point for calculating standard time values. See also *SNTP*.

UTP cable—Unshielded Twisted Pair cable. UTP cable consists of two wires twisted two or more times per inch in order to help cancel out noise. The entire cable has no covering. UTP cable is typically used in telephone lines for voice service, ARCnet networks, 10BaseT Ethernet networks, and particular sections of Token Ring networks. See also *10Base-T*.

UTP Ethernet—See *10Base-T*.

UTP port—Unshielded Twisted Pair port. The UTP port is an Ethernet port on the shelf-controller backpanel. Using the UTP port, you can connect a 10Base-T cable to the Pipeline 220. See also *UTP cable*.

V.24—An ITU-T standard that specifies a Physical-layer interface between Data Terminal Equipment (DTE) and Data Circuit-Terminating Equipment (DCE). V.24 is nearly identical to RS-232. See also *DCE*, *DTE*, *RS-232*.

V.35—An ITU-T standard for high-speed synchronous data transmission and exchange. In the U.S., most routers and Data Service Units (DSUs) that connect to T1 lines use V.35. See also *DSU*, *router*, *synchronous transmission*, *T1 line*.

Van Jacobson compression—See *VJ compression*.

Variable-Length Subnet Mask—See *VLSM*.

Virtual Private Network—See *VPN*.

VJ compression—Van Jacobson compression. VJ compression is a method for compressing Transmission Control Protocol (TCP) headers in order to decrease round-trip times on Serial Line Internet Protocol (SLIP) connections. The version of SLIP implementing VJ compression is called Compressed Serial Line Internet Protocol (CSLIP). See also *compression*, *CSU*.

VLSM—Variable-Length Subnet Mask. A VLSM is a way to configure an IP subnet for maximum flexibility. Two different subnets of the same IP network number may have different masks and, therefore, different sizes. A packet is routed to the longest or most specific match. VLSM is also referred to as *Classless Inter-Domain Routing (CIDR)*. See also *IP subnet*, *subnet mask*.

VPN—Virtual Private Network. A VPN is a private network that uses the Internet to carry all traffic. It can link all the offices, telecommuters, travelling employees, customers, and suppliers for a single organization. A VPN is virtual because it appears to the organization as a private network. All users see only their own traffic. See also *private network*.

VT-100—An ASCII-character data terminal, consisting of a screen and keyboard. Manufactured by Digital Equipment Corporation (DEC), the VT-100 has become an industry standard data terminal. VT-100 emulation software allows a standard PC to act as a VT-100 terminal.

WAN—Wide Area Network. A WAN is an internet of devices, generally consisting of several networks distributed over a wide geographic distance, connected by telephone lines, and using different hardware platforms and protocol encapsulation. See also *internet*.

WAN connection—A connection between two endpoints over a WAN, as opposed to a local connection by a serial or Ethernet link. See also *WAN*.

Wide Area Network—See *WAN*.

Windows Internet Name Service—See *WINS*.

WINS—Windows Internet Name Service. WINS is a Microsoft product that manages the mapping between resource names and IP addresses. The Domain Name System (DNS) service used on the Internet cannot dynamically map IP addresses to local resource names. Through dynamic database updates, WINS lets a user gain access to network resources by means of user-friendly names, rather than by means of IP addresses.

wiring hub—See *hub*.

X.75—The ITU-T international standard for connecting packet-switched networks. See also *packet switching*.

A

- Activation parameter 4-4
- address
 - MAC 12-2
- address (SNMP manager) 15-5
- Address or Filename field 3-3
- address security 15-1, 15-2
 - SNMP 15-1
- address spoofing 13-16
- adjacencies
 - forming 8-4
 - OSPF 8-5
- Advertise Dialout Routes parameter 6-11
- AEP 11-1
- Agent Mode 14-3, 14-10
- alarm events 15-8
- Allow only approved SNMP managers listed below 15-2
- ALU
 - described 16-12
 - displaying 16-12
- Always Spoof parameter 7-5
- AMI line encoding 4-2
- Answer tab 3-9
- AppleTalk
 - default zone 11-6
 - NBP Broadcast Request 11-4
 - necessary client software 2-3
 - seed vs non-seed 11-6
 - ZIP Query 11-4
 - zone multicasting 11-2
- AppleTalk Chooser 11-4
- AppleTalk Control Protocol (ATCP) 11-1
- AppleTalk Echo Protocol (AEP) 11-1
- AppleTalk networks (extended/non-extended) 11-2
- AppleTalk routing
 - configuring 11-5
 - how it works 11-4
 - per connection 11-7
 - RTMP packets 11-3
 - seed router 11-3
 - when to use 11-1
- AppleTalk zones 11-2
- area routing (OSPF) 8-6
- areas
 - specifying for OSPF 8-10
 - viewing OSPF 16-26
- ARP
 - and bridging 12-12
 - displaying cache 16-19
 - inverse 6-9
 - proxy 6-9
- ARP broadcasts 12-2
- ARP-directed (RFC 1433) x
- AS 8-2
 - OSPF 8-2
- ASBR 8-2
 - calculations 8-3
 - disabling calculations 8-11
- Ascend Configurator
 - disabling Syslog 3-6
 - features 1-7
 - installing 3-2
 - parameter naming convention 3-8
 - using 3-6
- Ascend Customer Service
 - contacting iii
 - international phone numbers iii
- Ascend Enterprise MIB 15-1, 15-8, 15-10
- Ascend Enterprise traps 15-8
- Ascend MIB 15-10
- ASE 8-2
- ASE tag 8-11
- ASE Tag parameter 6-28
- ASE type 8-11
- ASE Type parameter 6-28
- ATCP 11-1
- ATMP 1-3, 14-1
 - default route preference 6-5
 - foreign agent 14-1
 - gateway mode 14-9
 - gateway mode (IP) 14-10
 - gateway mode (IPX) 14-13
 - home agent 14-1
 - router and gateway mode 14-2
 - router mode (IP) 14-4
 - router mode (IPX) 14-8
- Attenuation (T1 line) 4-2
- authentication (supported protocols) 1-5
- authenticationFailure 15-9
- autonomous system 8-2
- Autonomous System Border Router, see ASBR.

B

- B8ZS line encoding 4-2
- backbone area 8-6
- backbone router 6-14
- backpanel, illustrated 2-4
- Backspace key 16-6
- Back-Tab key 16-6
- backup routers 8-4

- bandwidth
 - displaying for each session 16-11
 - frame relay 5-1
 - battery specifications A-1
 - BCP (RFC 1638) ix
 - Become Default Router parameter 7-5
 - black-hole interface 6-6
 - BOOTP 7-1
 - BOOTP relay 6-10
 - Bootstrap Protocol, see BOOTP
 - bridge table 12-2
 - managing 12-7
 - bridge-table 12-11
 - bridging
 - ARP broadcasts 12-2
 - broadcast addresses 12-2
 - configuring proxy mode 12-12
 - disadvantages 12-1
 - displaying number of sessions 16-10
 - enabling 12-3
 - example 12-10
 - in AppleTalk environment 11-2
 - introduction 1-6
 - most common uses 12-1
 - PPP-encapsulated link 12-5
 - static table entries 12-7
 - transparent/learning 12-7
 - when to use 12-1
 - broadcast addresses (and bridging) 12-2
 - broadcast IP address 6-3
 - BRs 8-4
 - Buildout parameter 4-2
 - when to add attenuation 4-4
- C**
- cabling pinouts A-3
 - Call Detail Reporting 6-37
 - CCP (RFC 1962) ix
 - CDR 6-37
 - channels
 - getting status information about 16-9
 - CHAP
 - RFC 1994 ix
 - Chooser 11-4
 - CIDR 8-3
 - RFC 1519 x
 - circuit (frame relay) 5-3, 5-13
 - circuit, displaying Frame Relay 16-34
 - class A default mask 6-2
 - class B subnet mask 6-2
 - class C subnet mask 6-2
 - classes of SNMP traps 15-5
 - Classless Inter-Domain Routing (CIDR)
 - RFC 1519 x
 - classless inter-domain routing, see CIDR.
 - client DNS 6-10
 - clients (displaying multicast) 16-24
 - Clock Source parameter 4-2
 - clock, maximum acceptable for V.35 4-4
 - CLU
 - described 16-12
 - displaying 16-12
 - coldStart 15-8
 - collisions, displaying number of Ethernet 16-13
 - command line described 16-16
 - commands
 - displaying terminal server 16-16
 - help for terminal server 16-17
 - Set 16-18
 - Show 16-18
 - Show Revision 16-34
 - Show Uptime 16-34
 - Show Users 16-34
 - terminal server 16-16
 - unsupported terminal server 16-17
 - communications software
 - necessary for serial communication 2-3
 - community strings 15-1
 - SNMP 15-2
 - Comparison 13-8
 - configuration
 - accessing interface using Telnet 16-2
 - accessing VT100 16-2
 - logging in as superuser 16-7
 - overview of VT100 interface 16-1
 - saving changes 16-5
 - setting permissions for 16-7
 - terminal emulation settings 16-2
 - uploading 3-13
 - using VT100 interface 16-3
 - configuration file
 - opening 3-13
 - saving 3-14
 - configuring circuit 5-13
 - configuring gateway connection 5-9
 - connected routes 6-5
 - connecting from VT100 interface 5-4
 - connections 16-32
 - displaying information about TCP 16-30
 - Connections tab 3-9
 - consoleStateChange 15-9
 - Control port

accessing interface using 16-2
connecting workstation to 16-2

cost
 IP routing 6-28
 OSPF 8-5, 8-11

CRC 16-13

D

D4 line framing 4-2

Data Communications Equipment, see DCE.

Data Link Connection Identifiers, see DLCI.

Datagram Delivery Protocol (DDP) 11-1

datalink 5-4, 5-13

DB-44 port 4-4

DBA, displaying whether installed 16-14

DCE 4-1

DCE Error Thresholds 5-5

DCE Event Count 5-5

DDP 11-1

default preference 6-5

default route
 and backbone router 6-14
 configuring 6-29
 Ignoring 6-34
 ignoring 6-9
 IPX RIP 10-2

default router (DHCP) 7-5

default subnet mask 6-2

default zone 11-6

Delay between messages 5-5

Delay to wait for messages before recording an error 5-5

Delete key 16-6

described and illustrated 2-9

designated routers 8-4

Destination address 13-8

destination field 6-4

DHCP 7-3

 configuring 7-4
 default router 7-5
 how addresses are assigned 7-4
 renewing addresses 7-4
 server 7-5
 spoofing 7-5

DHCP spoofing 7-4

Directed ARP (RFC 1433) x

displaying statistics about Frame Relay 16-32

DLCI 5-3

 inactive 5-4

DLCIs

displaying status 16-33
stopping traffic without disrupting 16-18

DNS 6-10

 client 6-10
 configuring 6-15
 creating local table 16-41
 deleting entry from local 16-43
 displaying entries 16-41
 displaying table 16-40
 editing local table 16-42
 monitoring table 16-41

DNS host table (local) 7-6

DNS lists 6-10

DNS query 6-10

DNS table (local) 7-7

document conventions viii

documentation set viii

domain name 6-10

Down-Arrow key 16-6

DRs 8-4

DS1 MIB implementation 15-10

dual IP 6-8

 example 6-8

dual lan access

 using the Pipeline 220 for 1-2

Dyn Stat status window 16-11

dynamic IP routes 6-4

dynamic routes 6-34

E

E1 line

 getting status information about 16-9

edit menus, described 16-1

EGP 8-2

Encapsulation Options screen

 Answer profile 3-15

Encoding 4-3

environmental requirements A-2

errors

 displaying errors during session 16-13

 Telnet messages 16-45

 viewing OSPF 16-25

ESF line framing 4-2

Ether Opt status window 16-14

Ether Stat status window 16-13

Ethernet

 connecting to 2-5

 creating IP interface 6-5

 displaying hardware type 16-15

 displaying statistics of 16-13

F

- enabling RIP 6-9
- necessity for supporting computer 2-2
- supplied cabling 2-1
- Ethernet interface
 - configuring OSPF 8-12
 - displaying MAC address of first 16-15
 - displaying MAC address of second 16-15
 - primary IP address 6-8
 - second IP address 6-8
 - specifications A-3
- events, getting system 16-9
- eventTableOverwrite 15-8
- Experience with the OSPF protocol
 - RFC 1246 x
- extended AppleTalk networks 11-2
- Exterior Gateway Protocol, see EGP.
- exterior routing protocols 8-2

F

- FCC Part 15 C-2
- FCC Part 68 Notice C-2
- features and overview 1-7
- Filter Action 13-8
- filter conditions 13-7
- Filter profile 13-6
- filters 1-5
 - activating 13-7
 - applying in Answer profile 13-2
 - applying in Connection profile 13-4
 - applying to Ethernet interface 13-4
 - call vs data 13-1
 - defining generic 13-8
 - defining IP 13-8
 - Destination address 13-8
 - example generic 13-9
 - example IP 13-16
 - forwarding action 13-1
 - input 13-7
 - input or output 13-6
 - IP 13-6
 - output 13-7
 - overview 13-6
 - Protocol 13-9
 - Source address 13-8
- Filters menu 13-6
- Filters tab 3-9
- Firewall-Friendly FTP (RFC 1579) x
- firewalls 1-5
 - displaying whether installed 16-14
- firewalls (RFC 1579) x
- flash RAM 1-8

- foreign agent 14-1, 14-2
- forwarding action 13-1
- Frame Relay
 - displaying circuit information 16-34
 - displaying DLCI status 16-33
 - displaying information about 16-32
 - displaying link management information 16-33
 - displaying statistics about 16-32
 - displaying whether installed 16-14
 - RFC 1586 x
 - turn off traffic without disabling endpoints 16-18
- frame relay 5-9, 5-13
 - bandwidth 5-1
 - circuit 5-3
 - circuits 5-9
 - configuring logical link 5-4
 - configuring NNI interface 5-5
 - configuring UNI-DCE interface 5-7
 - configuring UNI-DTE interface 5-8
 - connecting to switch 5-1
 - datalink 5-4, 5-13
 - DCE Error Thresholds 5-5
 - DCE Event Count 5-5
 - Delay between messages 5-5
 - Delay to wait for messages before recording an error 5-5
 - gateway connections 5-3, 5-9
 - link management protocol 5-4
 - logical interfaces 5-2
 - N391 5-5
 - N392 5-5
 - N393 5-5
 - nailed connection 5-4
 - NNI 5-2
 - NNI interface 5-15
 - Polling Cycles 5-5
 - T391 5-5
 - T392 5-5
 - UNI-DCE 5-2
 - UNI-DCE interface 5-13
 - UNI-DTE 5-2
- Frame Relay MIB implementation 15-10
- Frame Relay tab 3-9
- frames
 - displaying number of Ethernet received 16-13
 - displaying number of Ethernet transmitted 16-13
 - displaying number received on link 16-12
 - displaying number transmitted on link 16-13
- frDLCIStatusChange 15-8
- frequency and power requirements A-2
- ftp.ascend.com 15-10
- Full Access profile, permissions and 16-8

G

- gateway 5-3
- gateway connection 5-9
- gateway connections 5-9
- gateway field 6-4
- gateway mode (ATMP) 14-2
- General tab screen 3-4
 - Language selection 3-4
- generic filter 13-9
- generic filters 13-6
 - defining 13-8
- Generic Routing Encapsulation (GRE) 14-1
- GRE 14-1
- Greenwich Mean Time (SNTP) 6-11
- Guidelines for Running OSPF Over Frame Relay Networks
 - RFC 1586 x

H

- hardware-level address
 - and bridging 12-2
- HeartBeat Address 9-3
- Heartbeat Alarm Threshold 9-3
- HeartBeat Slot Time 9-3
- Hello packets 8-10
- HelloInterval 8-10
- History file radio button 3-6
- History file, naming 3-6
- home agent 14-1, 14-2
- Home Agent Password 14-3, 14-10
- Home Agent Type 14-3, 14-10
- hop count 6-28
- host addresses per class C subnet 6-3
- host table (DNS) 7-6

I

- IC CS-03 Notice C-3
- ICMP 6-4
 - RFC 1256 x
- ICMP packets, displaying 16-20
- ICMP Redirects 6-4
- ICMP redirects
 - and static routes 6-29
 - default preference 6-5
- ICMP Router Discovery Messages (RFC 1256) x
- ie0 interface 6-6

- IGMP 9-1
 - displaying statistics 16-23, 16-24
 - displaying statistics for groups 16-23
 - showing clients 16-24
- IGMP requests 9-3
- inactive DLCI 5-4
- inactive interface 6-6
- input filters 13-6, 13-7
- interface-based routing 6-6, 6-7
 - configuring 6-24
- interfaces
 - displaying statistics for 16-20
- Internet Assigned Numbers Authority (IANA) 7-12
- Internet gateway 1-3
- Internet Group Membership Protocol, see IGMP.
- Internet Protocol (IP) screen 3-15
- IP
 - displaying ARP cache 16-19
 - displaying ICMP packets 16-20
 - displaying routing table 16-35
 - displaying statistics and addresses 16-21
 - routing entries explained 16-35
 - routing table example 16-37
- IP address
 - assigning 3-10
 - broadcast address 6-3
 - default subnet mask 6-2
 - displaying 16-21
 - displaying remote device 16-12
 - verifying 6-24
- IP Address parameter 6-7, 6-8
- IP addresses
 - how DHCP assigns 7-4
 - in local DNS table 7-6
 - zero subnets 6-3
- IP filters 13-6
 - and IP/TCP/UDP 13-6
 - defining 13-8
 - example 13-16
- IP interfaces
 - Ethernet and internal 6-5
- IP Mobility (RFC 2002) ix
- IP route name 6-27
- IP routes
 - adding 16-40
 - black-hole, loopback, reject 6-6
 - default preferences 6-5
 - deleting 16-40
 - Ethernet interface 6-5
 - ie0 interface 6-6
 - inactive interface 6-6
 - metrics 6-5
 - multicast interface 6-6

- numbered interfaces 6-6
 - route preferences 6-5
 - Routes profile 6-4
 - WAN interfaces 6-6
 - IP routing
 - BOOTP relay 6-10
 - clients 6-19
 - configuring 6-27
 - configuring preferences 6-27
 - configuring remote address 6-17
 - configuring RIP policy 6-35
 - configuring Syslog 6-37
 - defining WAN interface 6-17
 - destination address 6-27
 - DHCP server 7-3
 - dual 6-8
 - dual IP example 6-8
 - dynamic route updates 6-34
 - ignoring default route 6-9, 6-34
 - introduction 1-6
 - inverse ARP 6-9
 - local domain name 6-10
 - metrics 6-18
 - name servers 6-10
 - poisoning routes 6-11
 - preferences 6-18
 - primary address 6-8
 - private routes 6-18, 6-28
 - proxy ARP 6-9
 - RIP policy 6-19
 - second address 6-8
 - UDP checksums 6-11
 - virtual hops and costs 6-28
 - WAN alias 6-18
 - IP routing table 6-4
 - at system startup 6-4
 - how Pipeline 220 uses 6-4
 - static and dynamic routes 6-4
 - IP static routes 6-29
 - IP Version 4 (RFC 1812) ix
 - IPCP (RFC 1332) ix
 - iproute show command 6-5
 - IPSec, displaying whether installed 16-14
 - IPX
 - checking NetWare servers 10-8
 - connecting Pipeline 220 to 10-8
 - displaying packet statistics 16-31
 - displaying routing table 16-32
 - displaying service table 16-31
 - IPXPing command 16-47
 - login.exe 10-3
 - Macintosh and UNIX clients 10-4
 - multiple frame types 10-1
 - necessary client software 2-3
 - packet burst 10-4
 - SAP 10-1
 - SAP broadcasts 10-1
 - static routes 10-11
 - viewing RIP/SAP tables 10-10
 - WAN considerations 10-3
 - IPX and AppleTalk Protocol screen 3-15
 - IPX filters 13-6
 - IPX network numbers 10-8
 - IPX preferred server 10-3
 - IPX RIP 10-2
 - configuring static route 10-22
 - default route 10-2
 - managing table 10-10
 - restricting 10-10
 - similarity to TCP/IP RIP 10-2
 - IPX RIP broadcasts 10-2
 - IPX Route profiles 10-2, 10-3
 - IPX routing
 - requirement of authentication 10-1
 - IPX SAP
 - filtering 10-15
 - managing table 10-10
 - restricting 10-14
 - IPX SAP filter
 - applying 10-6, 10-17
 - defining 10-15
 - IPX SAP filters 10-2, 10-3
 - IPXCP 10-1
 - IPXPing command, described 16-47
 - IPXWAN 10-1
 - ISDN, displaying whether signaling installed 16-14
- ## K
- keyboard commands, selection commands 16-6
 - keys
 - Backspace 16-6
 - Back-Tab 16-6
 - Delete 16-6
 - Down-Arrow 16-6
 - Left-Arrow 16-6
 - Tab 16-6
 - Up-Arrow 16-6
 - Kill command, described 16-49
- ## L
- Language selection
 - Ascend Configurator setup 3-4
 - learning bridge 12-7
 - leased line

- installation cautions 2-8
- LEDs 2-9
 - described 2-4
 - operation at startup 2-8
- Left-Arrow key 16-6
- Light Emitting Diode, see LEDs.
- Line status window 16-9
- lines, getting status information about 16-9
- link
 - getting status information about 16-9
 - quality of 16-11
- link state, displaying database 16-28
- Link this condition to the next... 13-8
- linkDown 15-8
- Link-State Advertisements, see LSAs.
- link-state routing algorithm 8-7
- linkUp 15-8
- local DNS table
 - creating 16-41
 - deleting 16-43
 - editing 16-42
- local domain name 6-10
- Local UDP port 14-3, 14-10
- Log tab 3-9
- logging
 - examples of Syslog 16-16
 - format of Syslog messages 16-15, 16-16
 - system events in status windows 16-9
- logical interfaces 5-2
- logical link 5-4
- login as superuser 16-7
- login.exe 10-3
- loopback interface 6-6
- LQM ix
- LSAs 8-4
 - displaying 16-29
 - displaying number created 16-27
 - displaying number received 16-27

M

- MAC address 12-2, 12-7
 - displaying 2nd Ethernet interface 16-15
 - displaying first Ethernet interface 16-15
 - displaying remote device 16-12
- Macintosh clients
 - as IPX clients 10-4
 - IP routing 6-19
- Management Information Base, see MIB.
- Match only established TCP connections 13-9
- Maximum Receive Unit, see MRU.
- maxTelnetAttempts 15-10
- MBONE 9-1
 - configuring interface 9-8
 - forwarding on a WAN link 9-8
- MBONE router 9-4
- Media Access Control, see MAC.
- menus
 - making active 16-3
 - opening 16-3
 - opening edit fields 16-4
 - terminal server 16-43
- message format (Syslog) 6-40
- messages (Syslog) 6-37
- metrics 6-5, 6-18
 - configurable OSPF 8-5
 - for learned RIP routes 6-28
- MIB 15-1
- MIB-II implementation 15-10
- modem cable 2-2
- Modem MIB implementation 15-10
- MP (RFC 1990) ix
- MPP, displaying errors during 16-13
- MRU 5-5
- Multicast
 - related RFCs x
 - RFC 1458 x
 - RFC 1584 x
- multicast
 - displaying activity 16-24
 - displaying clients 16-24
 - displaying statistics 16-25
 - displaying table for 16-23
 - IGMP 9-1
 - IP interface 6-6
- Multicast (RFC 1949) x
- multicast backbone, see MBONE.
- multicast clients 9-8
- Multicast Extensions to OSPF
 - RFC 1584 x
- multicast forwarding 9-3
- multicast heartbeat 9-2
- multicasting
 - configuring MBONE interface 9-8
 - configuring system-wide 9-4
 - configuring WAN interface 9-6
 - configuring WAN interfaces 9-9
 - MBONE router 9-4, 9-8
 - prioritized packet discarding 9-3
 - Rate Limit parameter 9-3
- multipath routing, described 16-38

N

- N391 5-5
- N392 5-5
- N393 5-5
- nailed connection 5-4
- Name Binding Protocol (NBP) 11-1
- Name Server (RFC 1877) ix
- name servers
 - DNS or WINS 6-10
- NBP 11-1
- NBP Broadcast Request 11-4
- NetWare
 - packet burst 10-4
 - WAN considerations 10-3
- NetWare servers
 - checking configuration 10-8
 - example configurations 10-19
- Network Address Translation (NAT) screen 3-16
- network numbers (IPX) 10-8
- Network tab screen 3-5
- Network-to-Network Interface, see NNI.
- new features
 - obtaining information iv
- NNI 5-2
- NNI interface
 - configuring 5-5, 5-15
- non-extended AppleTalk networks 11-2
- non-seed vs seed 11-6
- NSSA (RFC 1587) x
- numbered interface 6-24
- numbered interfaces 6-6

O

- Offset, Length, Mask, and Value 13-8
- On Internet Authentication (RFC 1704) x
- Open button 3-3
- OS2 clients 6-19
- OSPF 6-4
 - advantages over RIP 8-1
 - areas 8-6
 - AS 8-2
 - ASBR always on in Pipeline 16-27
 - ASBR calculations 8-3
 - ASE checksum 16-27
 - ASE type/ASE tag 8-11
 - autonomous system (AS) 8-2
 - configurable metrics 8-5
 - configuring 8-10
 - configuring on Ethernet 8-12

- configuring WAN 8-16
- cost 8-5
- disabling ASBR calculations 8-11
- displaying if ABR enabled 16-27
- displaying information about packets received and transmitted 16-30
- displaying interfaces 16-27
- displaying IP address of Pipeline Ethernet address 16-27
- displaying link-state database 16-28
- displaying LSAs 16-29
- displaying neighbors 16-29
- displaying number of external link-state databases 16-27
- displaying number of LSAs 16-27
- displaying number of LSAs received 16-27
- displaying routing table 16-29, 16-37
- displaying version 16-27
- displaying whether enabled 16-27
- DRs and BRs 8-4
- forming adjacencies 8-4
- HelloInterval 8-10
- how it adds RIP routes 16-39
- link-state 8-1
- link-state advertisements 8-4
- link-state routing algorithm 8-7
- monitoring 16-25
- RFC 1245 x
- RFC 1246 x
- route convergence 8-1
- security 8-3
- SFP algorithm 8-4
- third-party routing 16-39
- TOS support 16-27
- viewing area information 16-26
- viewing errors 16-25
- VLSM 8-3
- OSPF (RFC 1583) x
- OSPF costs
 - configuring 8-15
- OSPF intervals
 - configuring 8-15
- OSPF MIB (RFC 1850) x
- OSPF NSSA Option
 - RFC 1587 x
- OSPF protocol analysis
 - RFC 1245 x
- OSPF routes
 - default preference 6-5
- OSPF stub areas 8-6
- OSPF Version 2
 - RFC 1583 x
- OSPF Version 2 Management Information Base (RFC 1850) x
- output filters 13-6, 13-7

P

- packet burst 10-4
- packet filtering
 - related RFCs x
- packets
 - displaying ICMP 16-20
- parameters
 - changing enumerated 16-5
 - editing text in 16-4
- password (Telnet) 6-10
- passwords
 - default for Full Access Security profile 16-7
 - Telnet password
 - and accessing VT100 interface 16-7
- Permanent Virtual Circuit, See PVC.
- Permanent Virtual Circuit, see PVC.
- phase (power requirements) A-2
- physical address 12-7
- physical addresses and bridge table 12-2
- Ping command, described 16-46
- Pipeline security profiles 16-7
- Plug and Play
 - Windows 95 and Windows NT 7-3
- Point-to-Point Protocol
 - RFC 1661 ix
- Point-to-Point-Tunneling Protocol, see PPTP.
- poisoning IP routes 6-11
- Polling Cycles 5-5
- port state change events 15-8
- portAcrPending 15-9
- portCarrier 15-9
- portCollectDigits 15-9
- portConnected 15-9
- portDTENotReady 15-9
- portDualDelay 15-8
- portHaveSerial 15-9
- portInactive 15-8
- portLoopback 15-9
- portRinging 15-9
- ports
 - described 2-4
 - displaying UDP port statistics 16-23
- Ports tab 3-9
- portUseExceeded 15-10
- portWaiting 15-9
- portWaitSerial 15-9
- power requirements A-2
- PPP
 - displaying errors during 16-13
- IPXCP 10-1
- PPP (RFC 1661) ix
- PPP Bridging Control Protocol (RFC 1638) ix
- PPP Challenge Handshake Authentication Protocol (RFC 1994) ix
- PPP Compression Control Protocol (RFC 1962) ix
- PPP in HDLC-like Framing (RFC 1662) ix
- PPP Internet Protocol Control Protocol (RFC 1332) ix
- PPP Internet Protocol Control Protocol Extensions for Name Server Addresses (RFC 1877) ix
- PPP Link Quality Monitoring (RFC 1989) ix
- PPP Multilink Protocol (RFC 1990) ix
- PPP Stac LZS Compression Protocol (RFC 1974) ix
- PPP Vendor Extensions (RFC 2153) ix
- PPTP 14-1
 - default route preference 6-5
 - limitations of Pipeline 220 14-1
 - support for 1-3
- preferences 6-18
 - IP routing 6-28
 - routing and 16-39
 - static routes 6-29
- preferred servers
 - IPX 10-3
- primary IP address 6-8
- priority
 - DR and BDR 8-10
- private routes 6-18
 - RIP 6-28
- profiles
 - description of 16-4
 - displaying which is in use 16-14
 - opening 16-3
- promiscuous mode 12-3
- Protocol filters 13-9
- proxy ARP, inverse ARP 6-9
- proxy mode 12-12
- PVC 5-3, 5-9

Q

- Q.922 address 6-9
- Q.933 A 5-4
- Q.933 Annex A 5-4
- Quickstart
 - exiting 3-16
 - using 3-14
- Quickstart screens 3-15

R

- Rate Limit parameter 9-3
- Recent Display
 - Ascend Configurator setup 3-4
- reject interface 6-6
- remote device, displaying name, IP address and MAC address of 16-12
- remote IP address 6-17
- Remote Management screen 3-16
- Report of IAB Workshop (RFC 1636) x
- Request For Comments, see RFC.
- Requirements for IP Version 4 Routers (RFC 1812) ix
- Requirements for Multicast Protocols (RFC 1458) x
- retransmit interval 8-11
- RFC
 - listing of related ix
- RFC 1213 15-10
- RFC 1315 15-10
- RFC 1317 15-10
- RFC 1406 15-10
- RFC 1696 15-10
- RFC 1701 14-1
- RFCs
 - IP routing ix
 - OSPF x
 - PPP ix
 - RFC 1974 ix
- RIP 6-4, 6-34
 - configuring IPX static route 10-22
 - configuring on WAN link 6-36
 - default route preference 6-5
 - disadvantages over OSPF 8-1
 - distance-vector metrics 8-1
 - enabling on Ethernet 6-9
 - hop count limit 8-1
 - how OSPF adds 16-39
 - IPX 10-2
 - IPX broadcasts 10-2
 - private routes 6-18
 - restricting 10-10
 - route convergence 8-1
 - tagging routes 6-28
- RIP broadcast
 - restricting 10-10
 - updates 6-4
- RIP policy 6-34
 - configuring 6-35
 - on WAN interface 6-19
- RIP summary 6-34
- RIP tables
 - managing 10-10
 - viewing 10-10
- RIP version 2, see RIP-v2.
- RIP-v1 6-19, 6-34
 - RIP policy and RIP summary 6-34
- RIP-v2 6-19, 6-34
 - effect of RIP policy/summary 6-34
 - recommendations 6-9
- RJ-48C 2-1
- route convergence
 - RIP vs OSPF 8-1
- route flooding
 - preventing 8-7
- route name
 - IP 6-27
- route preferences 6-5
 - example configuration 6-33
- router mode (ATMP) 14-2
- routes
 - adding IP routes 16-40
 - deleting IP routes 16-40
 - displaying IP routing table 16-35
 - how OSPF adds RIP 16-39
 - preferences 16-39
- Routes and Bridges tab 3-9
- Routes profile 6-4
- routing
 - adding IP routes to table 16-40
 - AppleTalk 11-4
 - AppleTalk per-connection 11-7
 - configuring AppleTalk 11-5
 - deleting IP routes to table 16-40
 - displaying ARP cache 16-19
 - displaying IP statistics and addresses 16-21
 - displaying IPX table 16-32
 - displaying multicast forwarding table 16-23
 - displaying multicast statistics 16-23, 16-25
 - displaying number of active sessions 16-10
 - displaying OSPF table 16-29
 - displaying table 16-35
 - IP routing table entries explained 16-35
 - IP routing table example 16-37
 - monitoring OSPF 16-25
 - multipath 16-38
 - OSPF and RIP 16-39
 - OSPF third party 16-39
 - preferences 16-39
 - viewing OSPF area information 16-26
- Routing in a Multi-provider Internet (RFC 1787) x
- Routing Information Protocol, see RIP.
- routing protocols
 - exterior 8-2
- Routing Table Maintenance Protocol (RTMP) 11-1
- RS232 MIB implementation 15-10
- RS-449 cable specifications A-5

RS-449/422 A-4
 RTMP 11-1
 RTMP packets 11-3

S

safety
 grounding iv
 instructions iv

SAM 6-40

SAP 10-1
 broadcast packets 10-1
 filtering 10-15
 restricting 10-14

SAP filter
 applying 10-6, 10-17
 defining 10-15

SAP filters 10-3

SAP Reply 14-3, 14-10

SAP tables
 managing 10-10
 viewing 10-10

Scalable Multicast Key Distribution (RFC 1949) x

screens
 General tab 3-4
 Network tab 3-5
 Syslog tab 3-6

second IP address 6-8

Secure Access Firewall (SAM) 6-40

Secure Access, displaying whether installed 16-14

Secure Operation of the Internet (RFC 1281) x

security
 default password for Security profile 16-7
 displaying whether IPSec is installed 16-14
 issues with Configurator 3-14
 OSPF 8-3
 permissions in Full Access Security profile 16-8
 related RFCs x
 restricting configuration permissions in default Security profile 16-7
 RFC 1245 x
 Security profile activated after power-on 16-7
 setting configuration permissions 16-7
 supported features 1-5

Security Considerations for IP Fragment Filtering
 RFC 1858 x

security events
 SNMP traps 15-9

Security profiles
 activated after power-on 16-7
 activating to access VT100 configuration interface 16-7

 default password for Full Access 16-7
 displaying which is in use 16-14
 logging into Full Access profile 16-7
 overview of 16-7
 permissions in Full Access 16-8
 restricting permissions in 16-7

Security tab 3-9

seed router 11-3
 vs non-seed 11-6

serial communication
 connecting to IBM-compatible 2-6
 connecting to Macintosh-compatible 2-7
 connecting using UNIX 2-7

serial number, displaying Pipeline 16-14

serial port
 computer access 2-2
 connecting to 2-6

serial WAN cabling specifications A-4

serial WAN data flow 4-4

serial WAN port 4-4

serial-cable adapter 2-1

Service Advertising Protocol, see SAP.

session
 characters in Status window 16-10
 displaying bandwidth of 16-11
 displaying uptime 16-11
 initiating Telnet 16-43
 quality of 16-11

session ID 16-35

sessions
 displaying active 16-34
 displaying error during PPP or MPP 16-13
 displaying ID 16-35
 displaying number of active 16-10
 initiating TCP 16-45
 Telnet commands 16-44

Sessions status window 16-10

Set command 16-18

Show command 16-18

show if totals command, described 16-21

Show Revision command 16-34

Show Uptime command 16-34

Show Users command 16-34

signaling
 displaying whether ISDN installed 16-14

signals
 controlling serial WAN data flow 4-4

Simple Network Management Protocol, see SNMP.

Simple Network Time Protocol (SNTP) (RFC 2030) ix

Simple Network Time Protocol, see SNTP.

Site Security Handbook (RFC 1244) x

SNMP 15-1, 15-2

- address security 15-2
- alarm trap and multicasting 9-2
- Ascend Enterprise traps 15-8
- community strings 15-1, 15-2
- default community string 3-1
- introduction 1-8
- MIB 15-1
- SNMP alarm events 15-8
- SNMP manager 15-1
- SNMP manager address 15-5
- SNMP Trap profiles 15-5
- SNMP traps 15-1, 15-5
 - classes 15-5
 - port state change events 15-8
 - security events 15-9
 - setting 15-5
- SNTP 6-11
 - RFC 2030 ix
- SNTP server addresses
 - specifying 6-11
- software
 - displaying loaded version 16-14
 - displaying version and load 16-34
- software upgrades 1-8
- Source address 13-8
- specifications
 - battery A-1
 - Ethernet interface A-3
 - RS-449 cable A-5
 - serial WAN cabling A-4
 - V.35 cable A-4
- SPF algorithm 8-4
- spoofing 13-16
- Stac LZS compression (RFC 1974) ix
- Startup screen settings 3-3
- static bridge-table entry 12-11
- static entries
 - bridge table 12-7
- static IP routes 6-4
- static route
 - configuring 10-22
 - defining to remote subnet 6-31
 - IPX 10-10
 - preferences 6-29
- Static routes 6-5
- static routes
 - IPX 10-11
 - IPX RIP 10-3
- statistics, displaying interface 16-20
- status windows
 - characters in Session Status 16-10
 - description of 16-1
 - Dyn Stat 16-11
 - Ether Opt 16-14
 - Ether Stat 16-13
 - introduction 1-8
 - Line status 16-9
 - making active 16-3
 - overview of 16-8
 - Sessions 16-10
 - Sys Options 16-13
 - System Events 16-9
 - WAN Stat 16-12
- stub area 8-6
- subnet address format for class C 6-3
- subnet mask 6-2
- subnets (zero) 6-3
- support for 1-3
- switched connections, displaying whether Pipeline can establish 16-14
- Sys Options status window 16-13
- Syslog
 - configuring 6-37
 - described 16-15
 - disabling on Configurator 3-6
 - examples of 16-16
 - format of messages 16-15
 - format of notice messages 16-16
 - messages 6-40
- Syslog daemon 6-38
- Syslog selection 3-6
- Syslog tab screen 3-6
- system
 - displaying software loaded 16-14
 - displaying type of 16-14
 - displaying uptime 16-14, 16-34
 - getting status information about 16-9
- system administration, overview of 16-1
- System Events Status window 16-9
- System Information screen 3-15
- system startup
 - building IP routing table 6-4
- system-based routing 6-7
- system-based vs interface-based routing 6-6
- systemUseExceeded 15-10

T

- T1 connection
 - AMI and B8Zs 4-2
 - attenuation 4-2
 - configuring 4-1
 - D4 and ESF 4-2
 - framing and encoding 4-2
 - introduction 4-1

- T1 line
 - clocking 4-2
 - configuring 4-2
 - encoding 4-3
 - getting status information about 16-9

- T1.617 Annex D 5-4

- T1.617D 5-4

- T391 5-5

- T392 5-5

- Tab key 16-6

- tabs

- screen size limitations 3-7

- tagging RIP routes 6-28

- TCP

- displaying statistics about 16-30
 - initiating session to remote host 16-45
 - well-known ports 7-12

- TCP connection, retry timeout 7-8

- TCP Timeout parameter 7-8

- TCP/IP

- BOOTP relay 7-1
 - necessary client software 2-3

- TCP/IP client software 6-19

- TCP/IP Plug and Play configuration 7-3

- Telnet

- accessing configuration using 16-2
 - command options 16-44
 - error messages 16-45
 - initiating session 16-43
 - reaching Pipeline using 16-2
 - session commands 16-44

- Telnet password 6-10

- temperature

- recommended ambient iv

- terminal server 16-16

- commands for users 16-18
 - commands not supported 16-17
 - configuring menu for users 16-43
 - displaying commands 16-16
 - exiting 16-17
 - getting help for 16-17
 - overview of interface 16-1
 - Set command 16-18
 - Show command 16-18
 - Telnet session 16-43
 - ways to access 16-16

- terminal settings for VT100 emulation 16-2

- terminal-server security 1-5

- TFTP Port selection 3-5

- timeout

- TCP connection retry 7-8

- topological database 8-4

- Traceroute (RFC 1393) x

- Traceroute command, described 16-48

- Traceroute Using an IP Option (RFC 1393) x

- Traceroute, example of 16-38

- transit delay 8-11

- transparent bridging 12-7

- Trap profiles 15-5

- traps

- Ascend Enterprise 15-8
 - classes 15-5
 - port state change events 15-8
 - security events 15-9
 - setting 15-5
 - SNMP 15-5

- type-1/type-2 metric 6-28

U

- UDP

- displaying listen table 16-23
 - displaying statistics 16-22
 - well-known ports 7-12

- UDP checksums 6-11

- UNI-DCE 5-2, 5-13

- UNI-DCE interface
 - configuring 5-7, 5-13

- UNI-DTE 5-2

- UNI-DTE interface 5-8

- Universal Time Configuration
 - SNTP 6-11

- UNIX clients

- as IPX clients 10-4
 - IP routing 6-19

- UNIX host and Syslog daemon 6-38

- Up-Arrow key 16-6

- uploading configuration 3-13

- uptime

- displaying Pipeline 16-34
 - displaying system 16-14

- User Interface

- Ascend Configurator setup 3-4

- users, displaying 16-34

- User-to-Network Interface, see UNI-DCE, UNI-DTE.

- using Quickstart 3-14

- using the Pipeline 220 as 1-3

V

- V.35 A-4

- V.35 cable
 - specifications A-4

- V.35 port
 - configuring 4-4
 - introduction 4-1
- V.35/RS-449 4-4
- variable length subnet masks, see VLSM.
- virtual hop count 6-28
- virtual hops 6-28
- Virtual Private Networking 1-3
- Virtual Private Networking, See VPN.
- Virtual Private Networking, see VPN.
- VLSM 8-3
- voltage
 - power requirements A-2
- VPN 1-3, 14-1
 - ATMP 14-1
 - introduction 1-6
- VT100 interface
 - accessing 16-2
 - activating menus/status windows 16-3
 - Ascend Configurator equivalents 3-9
 - changing parameters in menus 16-5
 - connecting workstation to Control port 16-2
 - DO DIAL command 5-4
 - DO HANGUP command 5-4
 - navigating 16-3
 - opening edit fields 16-4
 - opening menus and profiles 16-3
 - overview of 16-1
 - saving changes 16-5
 - settings for terminal emulation 16-2
 - special characters 16-5
 - Telnet password required 16-7

W

- WAN alias 6-18
- WAN connections 6-17
- WAN interface (multicasting) 9-6
- WAN Interface Configuration screen 3-16
- WAN interfaces
 - IP routing 6-6
 - supported 4-1
- WAN link (RIP) 6-36
- WAN Stat status window 16-12
- WAN statistics 16-12
- warmStart 15-8
- weight
 - IP routing 6-28
 - Pipeline 220 unit iv
- well-known TCP and UDP ports 7-12
- Windows 95 (Plug and Play) 7-3

- Windows clients (IP routing) 6-19
- Windows NT (Plug and Play) 7-3
- WINS 6-10
- WWW application
 - using the Pipeline 220 for 1-2

Z

- zero subnets 6-3
- ZIP 11-1
- ZIP Query 11-4
- Zone Information Protocol (ZIP) 11-1
- zone multicasting 11-2
- zone names and case insensitivity 11-2
- zones 11-4
 - AppleTalk 11-2
 - default AppleTalk 11-6