Pipeline 50, 75, and 130 Addendum

Includes the Pipeline 85

Ascend Communications

December 7, 1998 Part number 7820-0331-001

Contents

Configuration

1

Overview

The following new features might affect the way you configure your unit:

Read this first

New Pipeline 50 and 75 units have different back panels from older models, and can only use software version 5.0B or newer. The U versions are available now, and S/T versions will ship fourth quarter of 1997.

Warning: Do not use an older version of software with a new Pipeline 50 or 75. If you use an older version of software with a Pipeline 50 or 75, the unit will not function and you will need to return it to Ascend for replacement. You can disable an old or new unit if you download an older version of software.

Older Pipeline 75 has a switch

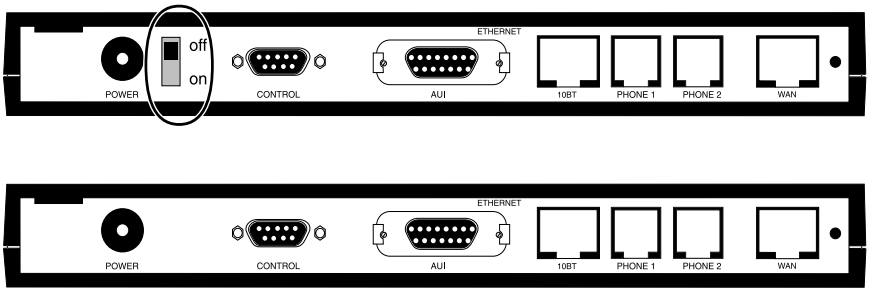

Newer Pipeline 75 does not have a switch

If you have a combination of old and new Pipeline 50 or 75 units, use these illustrations to determine which units are new. You can tell the difference between an older and newer units by looking for a switch on the back panel.

Older Pipeline 50 has a switch

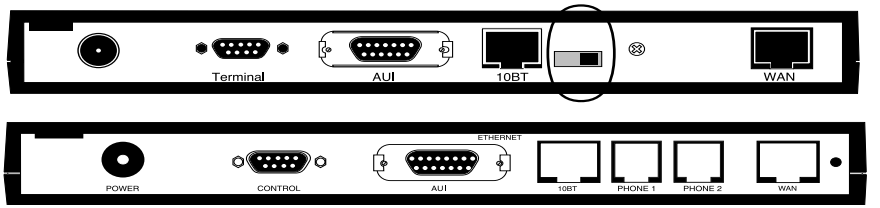

Newer Pipeline 50 does not have a switch

The Pipeline 50 and 75 units now have the same back panel.

Pipeline 85 installation and use

Installing the Pipeline 85

Before installing the Pipeline 85, take a look at the unit's back panel and familiarize yourself with the ports and labels, as shown below:

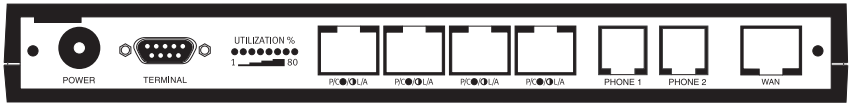

The ports and other elements on the back of the unit are, starting at the left:

- Power. Use the provided power supply to connect the unit to the current. Always plug the cable into the Pipeline before connecting it to the electrical current.

  Warning: Plugging the power supply into the wall socket before plugging the power cable into the Pipeline can create sparks, cause an electrical fire, or destroy the Pipeline.

  Because the Pipeline has no Power switch, plugging in the power supply turns the Pipeline on. After you plug in the Pipeline, it takes about a minute before it is ready to use.

- Terminal (serial) port. You can create a serial connect between your computer and the Pipeline by connecting a serial cable to a COM port on your computer, and connecting the other end to the terminal port on the Pipeline.
- Utilization percent lights. These indicate the percentage of LAN utilization on the Pipeline hub (between hosts on the local network). The percentage does not have anything to do with the traffic over the WAN connection.
- Four 10Base-T Ethernet ports. A straight-through cable from any host on the local network to the Pipeline connects the host to the Pipeline hub, thus expanding the network. You can connect the Pipeline hub to another hub on your LAN to make the Pipeline available to a greater number of users.

  The notation under the 10Base-T ports have the following meaning:

  – P/C means Partition/Collision
  – L/A means Link/Activity

  - The full circle indicates that when the LED is on, the partition or link on the Ethernet LAN is active. It does not apply to the WAN connection.
  - The half circle indicates that when the LED flashes, a collision or other activity is present on the Ethernet LAN, not the WAN connection.
- Phone ports 1 and 2. These are analog ports to which you can attach any type of phone equipment, such as phones, fax machines, caller ID device, or a modem. Please note that digital phone equipment, such as an ISDN telephone is not supported by the analog ports. The phone ports function identically to those on the Pipeline 75.
- WAN port is where the ISDN service is connected.

Please refer to the installation instructions for the Pipeline 75 for more detailed information. The main difference is the cabling of the Ethernet hub. Be sure you use a straight-through cable to connect to the hub. If you need to connect the Pipeline to another hub, refer to the documentation of the other hub for cabling information required to connect hub-to-hub.

Using the Pipeline 85 is identical to using the Pipeline 75

You can assume that any feature for a Pipeline 75 also applies to the Pipeline 85.

The Pipeline 85 functions identically to the version 2 Pipeline 75. Currently the version 2 Pipeline 75 includes all the features of the older Pipeline models, plus it has these additional features:

- On the Configure menu, the Caller ID and Forward disconnect parameters are available.
  - Caller ID set to No blocks your ISDN phone number on outgoing phone calls made from either analog Phone ports.
  - Forward disconnect initiates an off-hook click when the far end hangs up. It helps tear down a call when the Pipeline is behind a PBX.
- Multiple-address Network Address Translation (NAT) can be set. For all other models, only single-address NAT is available.
- Using the Pipeline as a hub does not require any software configuration on the Pipeline.

Pipelines can accept or reject voice calls

A new parameter, Ans Voice Call, enables you to configure a Pipeline unit to accept or reject voice calls.

The Pipeline 25 and 75 always accepted voice calls and directed them to the POTS ports, so you could not dedicate the unit to data calls. By setting Ans Voice Call to No, you can reject all incoming calls, and save the calling numbers that are rejected. Outgoing voice calls can still be made if a phone is attached to one of the Phone ports on the back of the unit.

The Pipeline 50 and 130 always considered analog calls to be Data Over Voice, and automatically accept the call. You can set the Pipeline 50 and 130 to reject these calls, as well, and the number of the rejected call is saved.

The Ans Voice Call parameter is in the Configure menu.

```
Configure...
   Switch Type=
   Chan Usage=
   My Num A=
   My Num B=
   SPID 1=
   SPID 2=
  >Ans Voice Call=Yes
```

The default is Yes, which accepts voice calls. If you set Ans Voice Call to No, you can list the phone number of rejected calls by entering the Show ISDN command at the terminal server prompt. For example:

```
ascend% show isdn

NL: CALL REJECTED/OTHER DEST: 5551010
```

The phone number listed, 5551010, originated a call that was not answered by the Pipeline.

Ans Voice Call

Description: Enables or disables incoming voice calls to a Pipeline with ISDN.

Usage: Enter Yes to enable or No to disable incoming voice calls.

Dependencies: None.

Parameter Location: Configure menu.

Configuration with a touch-tone telephone

For the Pipeline 75 only. This feature lets you use a touch-tone telephone to enter a Pipeline unit's ISDN telephone numbers. With this and two other features (automatic switch identification and remote configuration by a central configuration center), a person installing a Pipeline unit can quickly configure it without connecting it to a computer.

Overview

This procedure is intended primarily for telephone installers or other professional installers. It lets you quickly configure a Pipeline unit at the same time you install it. You do not need to connect a computer to the Pipeline unit to perform the procedure, but you do need a central configuration center that can call the Pipeline unit and finish the configuration after the Pipeline unit has established a connection to the ISDN line.

Note: This feature does not work on Pipeline units connected to Japanese NTT INS-64 telephone switches.

Required equipment and information

To configure a Pipeline unit with a touch-tone telephone, you need the following:

- A new Pipeline 75 unit with software version 5.0B or later.

  The procedure described in this note works for Pipeline units with factory default settings. To use the procedure with a Pipeline unit that has already been configured, install software version 5.0B or later if it is not already installed, open the terminal server, and then enter the `nvram` and `fclear` commands to restore the factory default settings.
- A cable for connecting the Pipeline unit to the ISDN telephone line.

  This cable (part number 2510-0122-001) is included with a new Pipeline unit.

- A touch-tone telephone and a modular telephone cable to connect it to the Pipeline unit.
- The telephone numbers of the ISDN telephone line that will be used by the Pipeline unit.
- The telephone number of the central configuration center that will complete the configuration.

Installing the Pipeline unit

To install the Pipeline unit:

1. Connect the Pipeline to the ISDN telephone line using the cable included in the box.
2. Connect a touch-tone telephone to the Phone 1 or Phone 2 jack on the back of the unit.

Configuring the Pipeline unit

To enter the telephone numbers for the ISDN telephone line:

1. Pick up the receiver of the touch-tone telephone.

   On all but a few older Pipeline units, you now hear an error tone. This tone is different from a busy signal.
2. Press the `*` key on the telephone.

   If you heard an error tone after step 1, the tone should now stop.
3. Press the `*` key on the telephone two more times.
4. Press the `1` key on the telephone.

   This specifies that you're entering the first of the telephone numbers for the ISDN line.
5. Enter the first telephone number for the ISDN line.
6. Press the `*` key on the telephone.

   Note: If you hear an error tone at any time after step 2, you have not entered the telephone number successfully. If this occurs, hang up the telephone and begin again with Step 1.

   If you've successfully entered the first telephone number, you now hear a busy signal. If so, enter the second telephone number:

7 Press the * key on the telephone.

8 Press the * key on the telephone two more times.

9 Press the 2 key on the telephone.

This specifies that you're entering the second of the telephone numbers for the ISDN line.

10 Enter the second telephone number for the ISDN line.

11 Press the * key on the telephone.

If you've successfully entered the second telephone number, you now hear a busy signal.

Once you've entered the ISDN telephone numbers, you let the Pipeline unit identify the type of telephone switch to which it's connected:

1 Hang up the touch-tone telephone.

The Pipeline unit identifies the switch type and sets its Switch Type setting to the appropriate value. If your ISDN telephone service uses service profile identifiers (SPIDs), the Pipeline unit also identifies these automatically and sets its SPID settings to the appropriate values. Automatic identification normally takes from one to three minutes.

2 Wait for the WAN status light to stop blinking.

When the light stops blinking, the automatic identification is complete. You can now make outgoing telephone calls with the telephone connected to the Pipeline unit. In addition, the Pipeline unit can now accept incoming calls.

To complete the Pipeline configuration:

1 Call the central configuration center.

You can call with the telephone connected to the Pipeline unit.

2 Give the central configuration center the ISDN telephone numbers for the Pipeline unit.

3 Have them call the Pipeline unit and enter the remaining configuration settings.

Automatic SPID detection - North America only

For the Pipeline 50 and 75 units installed in North America only. This feature adds a parameter that enables you to configure the Pipeline to automatically select the appropriate switch when the BRI is connected. If the switch is not AT&T P-to-P, the Pipeline will detect the SPIDs from the supplied phone numbers.

Configuring AutoSPID

To configure automatic switch selection on your Pipeline:

1 If you have an existing configuration (not the default configuration), save the configuration before proceeding.

2 Open the BRI... menu.

```
Configure
>Switch Type=AUTO SPID
 My Num A=
 My Num B=
 SPID A=N/A
 SPID B=N/A
 Data Usage=N/A
 Phone 1 Usage=N/A
 Phone 2 Usage=N/A
 Phone Num Binding=N/A
```

3 Select Switch Type=AUTO SPID.

4 Enter your full ten-digit phone numbers in My Num A and My Num B.

You must enter both phone numbers with the area codes. If these numbers are incorrect, this function will not be able to obtain the correct SPID.

The process may take several minutes. Do not enter the configuration again or unplug the BRI during this period.

After the switch type is detected, the link should fully initialize and display a D in the VT-100 Line Status window. If an M appears, the phone numbers may have been entered incorrectly, the Auto-SPID process has failed, or your BRI provider incorrectly provisioned your line. If this happens, enter

the phone numbers correctly or configure Switch Type and SPID1 and SPID2 manually.

Note: If you have arranged for AT&T point-to-point service, you do not have SPIDs associated with your service.

5 Save the configuration when prompted (after the switch type is properly detected). This will exit the BRI menu.

Data usage for additional switch types

For the Pipeline 75 only. The Data Usage parameter can now be set when the value of the Switch Type parameter is France, U.K., NET 3, Japan, Belgium, Australia, Swiss, German, or MP German.

Data Usage

Description: Specifies which of your ISDN telephone numbers to use for incoming data calls. If your ISDN service allows data calls on only one telephone number, you can use this parameter to specify the telephone number to use.

Usage: Press Enter to cycle through the choices.

- A allows incoming data calls to the telephone number specified by the My Num A parameter.
- B allows incoming data calls to the telephone number specified by the My Num B parameter.
- A + B allows incoming data calls to the telephone number specified by the My Num A parameter or the telephone number specified by the My Num B parameter.

Dependencies: If the value of the Switch Type parameter is AT&T/P-T-P, the Data Usage parameter is N/A. There is only one telephone number for this type of ISDN service, and this telephone number is used for all data calls.

Parameter Location: Configure Profile, BRI

See Also: My Num A, My Num B, Switch Type

Displaying the software load name

Ascend software releases are distributed in software *loads*, which are binary files that you copy to a local device and download to your Pipeline unit. Software loads vary according to functionality and target platform. The name of the software load is displayed in the Sys Options status window and in fatal error messages. The load name is an important aid to troubleshooting error conditions.

The name of the load file is in the format *function.model*. The prefix indicates the network interface and how the unit is designed to function. These abbreviations are some that are used (a complete listing is shown in the README file on the FTP server where the binary files are stored):

t	T1
e	E1
b	ISDN BRI
52	Switched 56 2 wire
54	Switched 56 4 wire
i	IP only
p	IPX only
x	X.25
1	old hardware (such as b1.p50)
2	new hardware (such as b2.p75)

For example:

b.p50 is for the Pipeline 50 BRI

52.p50 is for the Pipeline 50 Switched-56 2-wire

Pipeline models are abbreviated p50, p75, p85, and p13 for the Pipeline 50, 75, 85, and 130, respectively.

Note: When downloading the newest version of software from the Ascend FTP site (ftp.ascend.com/pub/Software-Releases), determine which file to download by referring to the README file associated with each sub-directory.

On your Pipeline unit, the current load appears in the Sys Options status window. For example:

```
00-100 Sys Option
>Access Router          ^
 Load: b.p75
 Switched Installed     v
```

Also see "TFTP checks compatibility of downloaded files" on page 7-24.

Alternate connection numbers for Pipeline 130

You can now add up to three alternate dial numbers in each connection profile of a Pipeline 130. The alternate numbers are tried before a designated secondary profile is called.

How alternate numbers are used to connect

Previously, each connection profile contained only one number to dial to reach the bridge, router, or node at the remote end of the link. If the call failed to connect using this number, information in a secondary profile (defined in the Session options) would be used to connect. Now if the dial number fails to connect, each of the alternate dial numbers is tried before the secondary profile is used.

Configuring alternate dial numbers

Three alternate numbers can be assigned to each connection profile. The alternate numbers are used in sequence from first to last, or until an alternate field is found to be empty (meaning if alternate number one and three are filled in, but number two is blank, the third number will never be attempted). If none of the alternate numbers reach the remote end of the link, the unit will use the designated secondary profile, if one is defined in the Session Options.

The Alternate Dial number parameters appear as follows:

```
Ethernet > Connections > profile
    Station=name
    Active=Yes
    Encaps=MPP
    Dial#=5551111
    Alt Dial#1=5551112
    Alt Dial#2=5551113
    Alt Dial#3=5551114
    Calling#=
    Route IP=Yes
    Bridge=No
    Dial Brdcast=N/A
    Encaps Options...
    IP Options...
    IPX Options...
    Session Options...
    Telco Options...
```

If you add two alternate numbers, be sure to use Alt Dial#1 and Alt Dial#2. A blank field indicates the end of alternate dial number information.

Alt Dial#n

Description: Up to three alternate dial numbers may be entered for the Pipeline 130 connection profiles. Use the alternate dial number fields in sequence. A blank field indicates the end of alternate numbers. (If Alt Dial#2 is blank, Alt Dial#3 will not be used, even if it contains data.)

Usage: Press Enter to open a text field and enter a phone number. Press Enter to close the text field. The field accepts up to 37 characters, which are limited to the following set:

```
1234567890()[]!z-*#|
```

Only the numerical characters are sent.

The default value is null.

Example: Alt Dial#1=5551112

Dependencies: An alternate number is used only if the number in Dial# does not connect, and only if a preceding alternate number field is not blank.

Parameter Location: Ethernet > Connections > profile.

72-character user name support

The My Name= parameter can now contain long names up to 72 characters.

Increased dial-out digits to 24

The maximum phone number length has been increased to 24 digits.

The following phone number fields are now 24 digits:

Configure profile > My Num A

Configure profile > My Num B

Ethernet > Connections > *any profile* > Dial #

Ethernet > Connections > *any profile* > Telco Options > Bill #

Wide Area Network Configuration 2

Overview

The following new features might affect the way you configure your unit:

Max Channel Count parameter changed

You can specify only supported values for the maximum number of channels to use for an MP+ call, set in the Max Ch Count parameter. For example, the Max Ch Cnt parameter on a Pipeline 50 can be set to a maximum of 2.

Previously, any value from 0 to 32 was allowed, regardless of whether the unit supported 32 channels or not.

The Max Ch Count parameter is located in Ethernet > Answer > *any profile* > PPP options, and Ethernet > Connections > *any profile* > Encaps options.

Interface support for MP call management

PPP connections are single-channel connections that connect to any other device running PPP. MP and MP+ are enhancements to PPP for supporting multi-channel links. In previous releases, if a connection was set up for "MPP," the Pipeline first requested MP+. If the other side of the connection didn't support MP+, the Pipeline would then request MP. If that protocol was also refused, PPP would be used instead.

You can explicitly configure the RFC 1717 MP option. MP supports multi-channel links, but not dynamic bandwidth allocation (DBA). The base-channel count is used to determine the number of calls to place, and the number of channels used for that connection does not change. In addition, MP requires that all channels in the connection share the same phone number (that is, the channels on the answering side of the connection must be in a hunt group).

These are the new parameters for configuring MP connections:

```
Ethernet > Connection > profile > Encaps=MP
Ethernet > Answer > Encaps > profile > MP=Yes
```

BACP support over MP added

Bandwidth Allocation Control Protocol (BACP) is the Internet standard equivalent to Ascend Multilink Protocol Plus (MP+).

How BACP works

BACP runs over MP and enables a unit from any vendor supporting MP to add or remove bandwidth on the basis of demand. BACP functions similarly to MP+ and uses the same menu items for MP+.

Since BACP does not support Idle Percent, the field has been removed from the Ethernet > Connection profile > Encaps menu.

Configuring BACP

To configure a Pipeline to use BACP for sending or receiving, set the BACP parameter in the Connection or Answer profile as follows:

- When sending, set Ethernet > Connection profile > Encaps > BACP=Yes (the default is No).

  Note: The Idle Percent field has been removed from this menu since it is always N/A for MP and is not supported by BACP.

- When receiving, set Ethernet > Answer > PPP BACP=Yes (the default is No).

This parameter is available only if Encaps=MP.

Pipeline MPP drops newest B channel first

The order in which MPP drops channels when bandwidth is decreased has changed. MPP drops the most recently connected channel first.

Inverse ARP for Frame Relay

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

Inverse ARP requests *must* be of the following type:

- ARP protocol type of IP (0x8000)

- ARP hardware address type is the 2-byte Q.922 address

All other types are discarded.

The Inverse ARP response supplies the following data:

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

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

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

Pipeline can be ATMP Home Agent

Virtual private networks can now include the Pipeline as a Home Agent ATMP end point. In this implementation, the Pipeline operates in router mode only.

Using a Pipeline in a virtual private network

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

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

Foreign and home agents

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

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

The home agent may communicate with the home network through a direct connection, another router, or across a nailed connection. When it relies on packet routing to reach the home network, it operates in router mode. It is in gateway mode when it has a nailed connection to the home network.

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

Configuring a home agent in router mode

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

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

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

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

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

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

```
Ethernet
  Connections
    Station=foreign-agent
    Active=Yes
    Encaps=MPP
    Dial #=555-1213
    Route IP=Yes
    Encaps options...
      Send Auth=CHAP
      Recv PW=foreign-pw
      Send PW=home-pw
    IP options...
      LAN Adrs=10.65.212.226/24
```

RIP-v2 can use Multicast IP or MAC address

Previously, the Pipeline could only broadcast RIP2 packets, which was compatible with RIP1 packet routing, but did not fully comply with RFC 1723, which recommends that there should be four possible methods, which are:

- Send onlyRIP1 packets.
- Broadcast RIP 2 packets.
- Multicast RIP 2 packets.
- Disable RIP.

A new parameter, RIP2 Use Multicast, has been added that enables Multicast RIP 2 packet support. The existing Ethernet > Mod Config > Ether options > RIP parameter support the other three modes.

RIP2 Use Multicast

Description: Specifies that Multicast IP is to be used for RIP 2 packets.

Usage: The possible values are Yes or No.

- No sends out RIP 2 packets using the settings of the Ethernet > Mod Config > Ether options > RIP parameter.

  No is the default.
- Yes sends RIP 2 packets with the IP Multicast address of 224.0.0.9 and the Multicast MAC address, and receives packets with a Multicast MAC address.

Dependencies: The RIP2 Use Multicast parameter does not apply if the Pipeline does not support IP (Route IP=No).

Parameter Location: Ethernet > Mod Config > Ether options

See Also: RIP

Pipeline 130 V.35 Serial WAN port

The Pipeline model P130-1UBRI-V35, provides a V.35 Serial WAN port in addition to an ISDN BRI or Switched 56 connection. The following sections describe how to configure and monitor the port.

Configuring the V.35 serial WAN port

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

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

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

- For Connection profiles, set the Group parameter in the Connections/Telco options submenu
- For a Frame Relay profile set the Nailed Grp parameter.

To configure the V.35 Serial WAN port:

1 Open the Serial WAN profile.

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

2 Open the Mod Config profile.

```
20-B00 Mod Config...
  Module Enabled-Yes
  Group=3
  Activation=Static
```

3 Select Module Enabled to Yes.

4 Select a Group number that matches the Group parameter in a Connection profile or the Nailed Grp parameter in a Frame Relay profile.

5 Select the appropriate Activation.

This selects the signals at the serial WAN port that indicate that the Data Circuit-Terminating Equipment (DCE) is ready to connect.

6 Close the Serial WAN profile and save the changes.

Parameter reference

This section describes the new parameters added to the Pipeline to support the V.35 serial WAN port.

Module Enabled

Description: Enables the Pipeline V.35 Serial WAN port.

Usage: Press Enter to cycle through the choices:

- Yes enables the Serial WAN port.
  Yes is the default.
- No disables the Serial WAN port.

Parameter Location: V.35 > Serial WAN > Mod Config

Activation

Description: This parameter selects the signals at the serial WAN port that indicate that the Data Circuit-Terminating Equipment (DCE) is ready to connect.

Flow control is always handled by the Clear To Send (CTS) signal.

Usage: Press Enter to cycle through the choices.

- Static specifies that the Pipeline does not use flow control signals because the DCE is always connected.
- DPR (Call Digit or Tone) specifies that the DCE raises the DPR signal when it is ready.
- CRQ (Call Request) specifies that the DCE raises the CRQ signal when it is ready.
- RTS (Request to Send) specifies that the DCE raises the RTS signal when it is ready.
- CRQ+RTS specifies that the DCE raises the CRQ and RTS signals when it is ready.
- DPR+CRQ+RTS specifies that the DCE raises the DPR, CRQ and RTS signals when it is ready.
- Disabled specifies that the V.35 serial WAN port is disabled. This setting will terminate an active session and prevent further attempts to establish a connection.
- Serial WAN profile: Serial WAN/Mod Config

Parameter Location: V.35 > Serial WAN > Mod Config

Group

Description: Assigns a group number to the serial WAN nailed channels. When the Group parameter of a Connection profile or the Nailed Grp parameter of a

Frame Relay profile have the same value as the Group parameter in the Serial WAN profile, the Connection or Frame Relay profile uses the serial WAN port.

Usage: Press Enter to open a text field. Enter a number between 1 and 60. The default is 3.

Press Enter again to close the text field.

Dependencies: Keep this additional information in mind:

- If you add channels to the Group parameter and save your changes, the Pipeline adds the additional channels to any online connection that uses the group.
- Do not assign more than one active Connection profile to a group.
- Do not assign a Connection profile to a group that a Frame Relay profile uses.

Parameter Location: V.35 > Serial WAN > Mod Config

See Also: Group, Nailed Grp

New Line Status field

The 10-100 Line Status window now contains a field next to the Link field that indicates the presence of the V.35 port and its status. The Pipeline monitors CTS and DTR signals to determine the link integrity.

If the V.35 port is present but inactive, a Line Status window similar to the following is displayed

```
10-100 1
Link   D   V.35
B1     *.....
B2     *.....
```

If the V.35 port is present and active, a Line Status window similar to the following is displayed

```
10-100 1
Link   D   V.35
B1     *....*
B2     *.....
```

Pipeline 130 V.35 physical specifications

The Pipeline weighs 1.25 pounds (563 g) and has these dimensions: 1.25" x 8.69" x 6.25" (3.28 cm x 22.1 cm x 15.9 cm). The Pipeline has the following physical interfaces:

Ports	Function or Operation
Power connector.	For 18 VDC power input.
Terminal (Control) port, type DE-9.	For system management and setup. 9600 bit/s (default), 8 bits per character, no parity bits, no flow control, and 1 stop bit.
One BRI port, type RJ-45, labeled WAN 1.	Far-right port: for BRI access to WAN. Factory options: U interface or S interface.
One V.35 port, type DB-44, labeled WAN 2.	Port adjacent to far-right: for V.35 access to WAN. See the Pipeline user documentation for channel usage restrictions.
Two Ethernet ports, type DA-15 (labeled AUI) and type RJ-45 (labeled UTP (10 BaseT)).	Single Ethernet interface, automatically selected by software when connected.

The table below shows the pinouts for the Pipeline unit's V.35 interface:

V.35 Signal	Female DB-44
FGND RI	1 8
SD+ SD-	39 40
RD+ RD-	30 29

V.35 Signal	Female DB-44
ST+ ST-	41 42
RT+ RT-	32 31
TT+ TT-	38 37
DTR DSR	6 11
DCD SGND	9 25
CTS RTS	7 36

Facilities Data Link (FDL) added for T1 lines

This feature provides full Channel Service Unit (CSU) capability to the Pipeline 130 T1.

How FDL works

A Facilities Data Link (FDL) is a 4 Kbps system-data channel available when using the T1 Extended Super Frame (ESF) format. FDL provides information at regular intervals to the carrier's maintenance devices enabling the telco to use the FDL protocol to monitor the quality and performance of T1 lines. Enabling FDL gives the Pipeline 130 T1 full CSU functionality.

This feature adds the FDL parameter to the Nailed T1 > Mod Config profile. The FDL parameter specifies the type of FDL protocol the Pipeline uses. Your carrier can tell you which FDL protocol to specify.

Keep this additional information in mind:

- The FDL parameter does not apply to D4-framed T1 lines.
- You continue to accumulate D4 and ESF performance statistics in the FDL Stats windows, even if you do not choose an FDL protocol.
  - D4 indicates the D4 format, also known as the Superframe format, for framing data at the physical layer.

    This format consists of 12 consecutive frames, separated by framing bits.
  - ESF indicates the Extended Superframe format for framing data at the physical layer.

    This format consists of 24 consecutive frames, separated by framing bits.

Configuring FDL

An FDL field has been added to the Nailed T1 > Mod Config menu.

To enable FDL:

1 Open the Nailed T1 profile.

2 Open the Mod Config menu.

```
90-B00 Mod Config
Nailed T1 Group=3
Activation=Enabled
Framing Mode=ESP
Encoding=B8ZS
>FDL=None
First DSO=1
Last DSO=6
Loop Back=Normal
```

3 Select the appropriate FDL mode. Press Enter to cycle through the choices. Your carrier will tell you which mode is available to you.

The options for the FDL field are:

Option	Description
ANSI	ANSI FDL implemented.

Option	Description
ATT	ATT FDL implemented.
Sprint	Sprint FDL implemented.
None	FDL not implemented. This is the default value.

Pipeline 130 T1 internal clock mode

You can now specify whether the Pipeline should generate the T1 transmit clock locally or synchronize to the network using the receive clock. Previously, the Pipeline did not generate the T1 transmit clock from its internal oscillator.

How Clock Source works

The Clock Source parameter is in the Nailed T1 > Mod Config submenu. The default is No, which is the typical setting used when connecting the Pipeline to a telco-provided T1 interface.

When connecting two Pipeline 130 units together on a private twin-shielded, twisted-pair connection, set Clock Source on one Pipeline to Yes, and set Clock Source on the other Pipeline to No.

Configuring the T1 transmit clock

To configure the Pipeline so that the T1 transmit clock is generated by the unit's internal oscillator, select Clock Source=Yes

```
30-100 Mod Config
   Nailed T1 Group=3
   Activation=Enabled
   Framing Mode=ESF
   Encoding=B8ZS
   First DSO=1
   Last DSO=24
  >Clock Source=No
```

```
Loop Back=Normal
```

Flexible start DS0 origin added for Nailed T1

For the Pipeline 130 only. Previously, the Nailed T1 configuration required that the first DS0 of the T1 line be the start of data. When adding a Pipeline 130 with Nailed T1 to an existing T1 line, this could cause a conflict when the first position was already used. You can now select a channel other than the first T1 channel as the origin for the first DS0.

How to set the DS0 origin

Configuring DS0 channels for Nailed T1

You specify the first and last DS0 channels in the Nailed T1 > Mod Config submenu. The number of channels you can specify depends upon the encoding mode you specify for the line. If the encoding mode is AMI, you cannot set more than six DS0 channels. The Last DS0 must be equal to (in the case of one channel) or greater than the first DS0.

To configure the DS0 channels for Nailed T1 on the Pipeline 130:

1. Open the Nailed T1 profile.

   If necessary press Escape until the main Edit menu is displayed, then select Nailed T1 and press Enter.

2. Select Mod Config.

   The Nailed T1 > Mod Config submenu appears. The values shown below are an example of values that might be entered in this menu

```
Mod Config
 Nailed Grp=1
 Activation=Enabled
 Framing Mode=ESF
 Encoding=B8ZS
>First DS0=1
 Last DS0=6
 Loop Back=Normal
```

3 Select the First DS0.

Set a start channel.

4 Select the Last DS0.

Set an ending channel. This number sets the upper limit of the range of DS0 channels and must be equal to or greater than the value in First DS0. For example, if you select First DS0=2 and Last DS0=7, you have indicated a range of six channels with 2 as the first channel. If your First DS0=2 and Last DS0=2, you have indicated that only one DS0 channel will be used (channel 2).

If the value of First DS0 is greater than the value of last DS0, this message appears:

```
Message #224: First DS0 value must be less than or equal
to Last DS0.
```

Note: If you select Alternative Mark Inversion (AMI) mode, do not set the total number of DS0 channels to be greater than six. B8ZS mode allows a maximum of 24 channels.

If you attempt to save the Nailed T1 profile with more than six channels, the following error message appears and you cannot save the profile:

```
Invalid channel count.
Maximum channel count may not exceed 6
```

5 Select Save, then press Enter to save your changes.

BACP/BAP PPP protocol IDs match IETF

The protocol IDs for BACP/BAP PPP on Ascend units have been changed to conform to the protocol IDs specified by the IETF. These changes are internal and cannot be modified by the user; there are no changes to the user interface.

Note: Both sides of a connection must use the same version of the BACP/BAP protocol IDs in order to connect successfully.

The protocol IDs that have changed are as follows:

Table 8.

Previous Protocol ID	Current Protocol ID
8071	c02b.
0071	c02d

IP Routing

3

Overview

The following new features might affect the way you set up IP routing on your unit:

IP Security

Note: To use this feature, you must obtain a hash code and download a special version of the code (IP-only load). Then you must have Secure Access Firewall and Secure Access Manager installed.

The Pipeline can encapsulate and decapsulate packets with encryption and authentication headers as described in RFCs 1826, 1827, 1828, 1829, 1851, and 1852. You can set up and regulate IP Security through the VT100 menus.

IP Security overview

The table below lists new IP security terms and concepts.

Concept	Description
Security Association (SA)	A collection of parameters and state variables used by a security transform for data traveling from one system to another. An SA is uniquely specified by a Security Parameters Index (SPI) and a destination IP address. It contains secrets for encryption, decryption, or authentication, as well as any other information necessary for the transform in use. Each SA describes only a single encapsulation for a single direction (such as outgoing ESP/ DES-CBC to a specific IP address).
Security Parameters Index (SPI)	A 32-bit integer value that the Pipeline uses, along with the destination address, to select an SA.
Security scheme	A template that includes the IP address of the remote end of a tunnel, four SAs and their associated SPIs, keys, and other transform-specific parameters. You set up this template using a set of new parameters in the VT100 interface. You need not configure all SAs, though you must configure at least one per active scheme. The Pipeline can only use the SAs you provide to encapsulate transmitted packets or to check received packets.

Concept	Description
Transform	A method of encapsulating the results of applying a security algorithm to a user packet, including any necessary keys, secrets, or parameters. An AH (IP Authentication Header, RFC 1826) authentication transform handles the process of authenticating users. The Pipeline supports AH-MD5 (RFC 1828) and AH-SHA (RFC 1852) authentication transforms. An ESP (IP Encapsulating Security Payload, RFC 1827) encryption transform handles the process of encrypting data. The Pipeline supports the ESP-DES-CBC (RFC 1829) and ESP-3DES-EDE-CBC (RFC 1851) encryption transforms.
Tunnel	You set up an IP Security tunnel with the local router as one endpoint, and the remote system as the other. The remote end of the tunnel is the system that verifies and decrypts the Pipeline's outbound security-encapsulated packets. After decapsulating the packet, the remote tunnel endpoint then passes the unencrypted and unauthenticated packet to its final destination.
Static scheme	A Security scheme with a fixed remote tunnel IP address.
Mobile scheme	A Security scheme without a configured remote tunnel IP address. The Pipeline learns the address of the remote tunnel endpoint from the first inbound packet in a session from the remote end. A mobile scheme is useful when the remote system is a customer of an ISP that dynamically assigns IP addresses to customers dialing in.
Scheme database SA database	At system startup, the Pipeline examines each scheme and enters it into the Scheme database. In addition, the Pipeline enters each of the four SAs into the SA database, along with parameters that you set. When you add, delete, or modify a scheme using the VT100 interface, the Pipeline updates the Scheme and SA databases to reflect the changes. Conflicts elicit an error message and defective schemes are not saved.

Setting up IP Security

To set up IP Security, follow the steps described below. For complete information on each parameter, see “New parameters” on page 3-7.

1 To enable the IP Security feature, specify the proper code for the IP Security parameter in the System > Feature Codes menu.

For example, your specification might look like this one:

```
20-400 Feature Codes
 >IP Security=ips-yg38-t22b-r+
```

2 Go to the Ethernet > IP Security menu, which contains a list of Security schemes you can configure.
For example:

```
20-600 IP Security
 >20-601 Harrison
  20-602 Jackson
  20-603 Anderson
  20-604
  20-605
```

3 Open a new scheme.
For example:

```
20-604
 >Name=
  Active=No
  Mobile=No
  Tunnel Address=0.0.0.0
  Received Auth...
  Transmitted Auth...
  Received Crypt...
  Transmitted Crypt...
```

4 For the Name parameter, specify the name of the scheme up to 16 characters in length. The default value is null.

5 Set Active=Yes.

6 For the Mobile parameter, specify whether the Security scheme has a fixed IP address for the remote endpoint of the tunnel.

If the Security scheme has a fixed IP address, set Mobile=No. If the Security scheme does not have a fixed IP address, set Mobile=Yes. The default value is No.

7 For the Tunnel Address parameter, specify the IP address of the remote system that verifies and decrypts the Pipeline's outbound security-encapsulated packets.

After decapsulating the packet, the remote tunnel endpoint passes the unencrypted and unauthenticated packet to its final destination. You set up an IP Security tunnel with the local router as one endpoint, and the remote system as the other.

8 To specify parameters for the authentication of received packets, open the Received Auth entry.
For example:

```
20-604 Easthampton
  Received Auth...
  >SPI=1
  Transform=None
  MD5 Key=N/A
  SHA-1 Key=N/A
```

Each value you specify in the Received Auth menu must match the values specified for the same parameters in the Transmitted Auth menu at the remote end of the IP Security tunnel.

9 For the SPI parameter, specify the Security Parameters Index (SPI) that the Pipeline uses, along with the destination address, to select a Security Association (SA).

Specify an integer between 1 and 4294967295. The default value is 1. The value you specify must be different from any other SPI configured on the router.

10 For the Transform parameter, specify an authentication transform.

MD5 indicates the AH-MD5 authentication transform (RFC 1828). SHA-1 indicates the AH-SHA-1 authentication transform (RFC 1852). None disables the Security Association (SA). The default value is None.

11 If Transform=MD5, set the MD5 Key parameter.

12 If Transform=SHA-1, set the SHA-1 Key parameter.

13 To specify parameters for the authentication of transmitted packets, go back to the Security scheme menu, and select the Transmitted Auth entry.
For example:

```
20-604 Easthampton
  Transmitted Auth...
  >SPI=1
  Transform=None
  MD5 Key=N/A
  SHA-1 Key=N/A.
```

14 Specify a value for the SPI and Transform parameters, and for the MD5 Key or SHA-1 Key parameter. Each value you specify in the Transmitted Auth menu must match the values specified for the same parameters in the Received Auth menu at the remote end of the IP Security tunnel.

15 To specify parameters for the encryption of received packets, go back to the Security scheme menu, and select the Received Crypt entry.
For example:

```
20-604 Easthampton
  Received Crypt...
  >SPI=1
  Transform=None
  DES Key=N/A
  DES IV Length=N/A
  3DES Key 1=N/A
  3DES Key 2=N/A
  3DES Key 3=N/A
  3DES IV Length=N/A
```

Each value you specify in the Received Crypt menu must match the values specified for the same parameters in the Transmitted Crypt menu at the remote end of the IP Security tunnel.

16 Set the SPI parameter.

17 For the Transform parameter, specify an encryption transform.

DES-CBC indicates the ESP-DES-CBC (RFC 1829) encryption transform. 3DES-CBC indicates the ESP-3DES-EDE-CBC (RFC 1851) encryption transform. None disables the Security Association (SA). The default value is None.

18 If Transform=DES-CBC, set the DES Key and DES IV Length parameters.

19 If Transform=3DES-CBC, set the 3DES Key and 3DES IV Length parameters.

20 To specify parameters for the encryption of transmitted packets, go back to the Security scheme menu, and select the Transmitted Crypt entry.
For example:

```
20-604 Easthampton
  Transmitted Crypt...
  >SPI=1
  Transform=None
  DES Key=N/A
  DES IV Length=N/A
  3DES Key 1=N/A
  3DES Key 2=N/A
  3DES Key 3=N/A
  3DES IV Length=N/A
```

21 Specify a value for each parameter.

Each value you specify in the Transmitted Crypt menu must match the values specified for the same parameters in the Received Crypt menu at the remote end of the IP Security tunnel.

22 Save your changes.

New parameters

3DES IV Length

Description: Specifies the number of bits in the initialization vector.

When the system uses a block cipher (like DES) in Cipher Block Chaining (CBC) mode, each cipher block is encrypted using the previous block. Because the first block of each packet does not have a previous block, the system includes an initialization vector with each packet in order to get the process started.

Usage: You can specify either 32 or 64 bits. The default value is 32.

Dependencies: For the 3DES IV Length parameter to apply, you must specify

the ESP-3DES-EDE-CBC encryption transform by setting Transform=3DES-CBC.

Parameter Location: Ethernet > IP Security > *scheme* > Received Crypt, and Ethernet > IP Security > *scheme* > Transmitted Crypt

See Also: Transform

3DES Key *n*

Description: Specifies the 3DES key, a secret shared between tunnel endpoints.

Usage: Specify a 64-bit hexadecimal value. Each byte must contain seven bits for the key and one bit (the 01 bit) for odd parity. If you enter and save a key with bad parity, a warning message displays and the Pipeline corrects the parity but does not save the profile. Saving a second time stores the profile.

Dependencies: For the 3DES Key parameter to apply, you must specify the ESP-3DES-EDE-CBC encryption transform by setting Transform=3DES-CBC.

Parameter Location: Ethernet > IP Security > *scheme* > Received Crypt, and Ethernet > IP Security > *scheme* > Transmitted Crypt

See Also: Transform

Active

Description: Specifies whether the IP Security scheme is enabled.

Usage: Specify one of these settings:

- Yes enables the scheme.
- No disables the scheme.

  The default value is No. When you choose this value, the name of the scheme appears with a hyphen and a space preceding it on the Ethernet > IP Security menu.

Parameter Location: Ethernet > IP Security > *scheme*

See Also: Mobile, Name, Tunnel Address

DES IV Length

Description: Specifies the number of bits in the initialization vector.

Usage: You can specify either 32 or 64 bits. The default value is 32.

Dependencies: For the DES IV Length parameter to apply, you must specify the ESP-DES-CBC encryption transform by setting Transform=DES-CBC.

Parameter Location: Ethernet > IP Security > *scheme* > Received Crypt, and Ethernet > IP Security > *scheme* > Transmitted Crypt

See Also: Transform

DES Key

Description: Specifies the DES key, a secret shared between tunnel endpoints.

Usage: Specify a 64-bit hexadecimal value. Each byte must contain seven bits for the key and one bit (the 01 bit) for odd parity. If you enter and save a key with bad parity, a warning message displays and the Pipeline corrects the parity but does not save the profile. Saving a second time stores the profile.

Dependencies: For the DES Key parameter to apply, you must specify the ESP-DES-CBC encryption transform by setting Transform=DES-CBC.

Parameter Location: Ethernet > IP Security > *scheme* > Received Crypt, and Ethernet > IP Security > *scheme* > Transmitted Crypt

See Also: DES IV Length, Transform

IP Security

Description: Specifies the IP Security feature code—a string of characters that enables the IP Security feature on the Pipeline. You purchase the feature code when you purchase the product. After you enter it for the IP Security parameter and reboot the Pipeline, the feature is enabled.

Usage: Specify the code for the IP Security feature. A correct code enables one of three levels of IP Security. You can tell the state of IP Security on the Pipeline

by looking at the Sys Option menu. Depending on the feature code you enter, one of these options appears:

- IPsec Not Inst—No IP security installed
- IPsec Inst 40 Bit—Security with up to 40-bit encryption keys installed
- IPsec Inst 56 Bit—Security with up to 56-bit encryption keys installed
- IPsec Unlimited—Security with arbitrary length encryption keys installed

Parameter Location: System > Feature Codes

MD5 Key

Description: Specifies the MD5 key, a secret shared between tunnel endpoints.

Usage: Specify a bit string with an even number of hexadecimal digits. You must specify at least two digits, but no more than 32. The default value is null.

Dependencies: For MD5 Key to apply, you must specify the AH-MD5 authentication transform by setting Transform=MD5.

Parameter Location: Ethernet > IP Security > *scheme* > Received Auth, and Ethernet > IP Security > *scheme* > Transmitted Auth

See Also: SHA-1 Key, SPI, Transform

Mobile

Description: Specifies whether the Security scheme has a fixed IP address for the remote endpoint of the tunnel. If it does not, the scheme is mobile.

The Pipeline does not initially enter a mobile scheme into the Scheme database. When the Pipeline receives a packet it can authenticate, the Pipeline marks the original mobile scheme busy. Then, it enters a copy of the scheme into the Scheme database, copying the tunnel address from the received packet into the newly created scheme. This procedure allows the Pipeline to use the newly created scheme to encapsulate transmitted reply packets, and to match further received packets from the same remote system.

A mobile scheme without an entry in the scheme database does not have an associated tunnel; therefore, the Pipeline cannot use it to encapsulate transmitted

packets. If a packet arrives and is authenticated in a way that matches a mobile scheme with a different IP address as the tunnel endpoint, the Pipeline changes the database entry to point to the new tunnel address. Therefore, the Pipeline can use a mobile scheme for only one remote system at a time.

Usage: Specify one of these settings:

- Yes indicates that the scheme is mobile.
- No indicates that the scheme is not mobile.
  The default value is No.

Dependencies: When Mobile=No, Tunnel Address does not apply.

Parameter Location: Ethernet > IP Security > *scheme*

See Also: Active, Name, Tunnel Address

Name

Description: Specifies the name of the IP Security scheme.

Usage: Specify a name containing up to 16 characters. The default value is null.

The Pipeline indexes each scheme by integer—1 (one) for the first scheme, 2 (two) for the second scheme, and so on. The Pipeline can dynamically allocate Security schemes, in which case it chooses the same number space, leaving adequate room for all possible static schemes.

Parameter Location: Ethernet > IP Security > *scheme*

See Also: Active, Mobile, Tunnel Address

SHA-1 Key

Description: Specifies the SHA-1 key, a secret shared between tunnel end-points.

Usage: Specify a bit string containing an even number of hexadecimal digits.You must specify at least two digits, but no more than 40. The default value is null.

Dependencies: For SHA-1 Key to apply, you must specify the AH-SHA-1 authentication transform by setting Transform=SHA-1.

Parameter Location: Ethernet > IP Security > *scheme* > Received Auth, and Ethernet > IP Security > *scheme* > Transmitted Auth

See Also: MD5 Key, SPI, Transform

SPI

Description: Specifies the Security Parameters Index (SPI) that the Pipeline uses, along with the destination address, to select a Security Association (SA).

Usage: Specify a decimal number between 1 and 4294967295. The default value is 1. The value you specify must be different from any other SPI configured on the router.

Parameter Location: Ethernet > IP Security > *scheme* > Received Auth,
Ethernet > IP Security > *scheme* > Transmitted Auth,
Ethernet > IP Security > *scheme* > Received Crypt, and
Ethernet > IP Security > *scheme* > Transmitted Crypt

See Also: MD5 Key, SHA-1 Key, Transform

Transform

Description: Specifies a method of encapsulating the results of applying a security algorithm to a user packet, including any necessary keys, secrets, or parameters.

Usage: In the Received Auth and Transmitted Auth menus, you can specify an authentication transform. Select one of these values:

- MD5 indicates the AH-MD5 authentication transform (RFC 1828).
- DES-CBC (40 Bit)
- SHA-1 (indicates the AH-SHA-1 authentication transform (RFC 1852).
- None disables the Security Association (SA).
  The default value is None.

In the Received Crypt and Transmitted Crypt menus, you can specify an encryption transform. Select one of these values:

- DES-CBC indicates the ESP-DES-CBC (RFC 1829) encryption transform.

- 3DES-CBC indicates the ESP-3DES-EDE-CBC (RFC 1851) encryption transform.
- None disables the Security Association (SA).

  The default value is None.

Dependencies: Keep this additional information in mind:

- When Transform=MD5, you must specify a value for the MD5 Key parameter. The SHA-1 Key parameter does not apply.
- When Transform=SHA-1, you must specify a value for the SHA-1 Key parameter. The MD5 Key parameter does not apply.
- When Transform=DES-CBC, you must set the DES Key parameter. The 3DES Key and 3DES IV Length parameters do not apply.
- When Transform=3DES-CBC, you must set at least one 3DES Key parameter. The DES Key and DES IV Length parameters do not apply.

Parameter Location: Ethernet > IP Security > *scheme* > Received Auth,
Ethernet > IP Security > *scheme* > Transmitted Auth,
Ethernet > IP Security > *scheme* > Received Crypt, and
Ethernet > IP Security > *scheme* > Received Auth

See Also: 3DES Key, 3DES IV Length, DES Key, DES IV Length, MD5 Key, SHA-1 Key

Tunnel Address

Description: Specifies the IP address of the remote system that verifies and decrypts the Pipeline's outbound security-encapsulated packets. After decapsulating the packet, the remote tunnel endpoint then passes the unencrypted and unauthenticated packet to its final destination.

You set up an IP Security tunnel with the local router as one endpoint, and the remote system as the other.

Usage: Specify an IP address in dotted decimal notation. The default value is 0.0.0.0.

Parameter Location: Ethernet > IP Security > *scheme*

See Also: Active, Mobile, Name

New diagnostic commands

This section describes the new diagnostic commands added to support the IP Security feature.

IPsecSADump

This command displays all SAs currently residing in the database. You enter the command without optional arguments:

>IPsecSADump

This type of information displays:

```
SA at 0x2c3f90 (scheme 1):
  flags=21 <ACTIVE,ESP,MOBILE>
  SPI=2, dst=Mobile(Unset), ESP type=DES3_CBC IV Size=64,
IV=0x26575d3ea7478374
SA at 0x2c3ed0 (scheme 1):
  flags=21 <ACTIVE,ESP,MOBILE>
  SPI=2, dst=Localhost, ESP type=DES_CBC IV Size=32,
IV=0x57f00cb6a80ff349
SA at 0x2c3e10 (scheme 1):
  flags=22 <ACTIVE,AH,MOBILE>
  SPI=1, dst=Mobile(Unset), AH type=SHA1
SA at 0x2c3d50 (scheme 1):
  flags=22 <ACTIVE,AH,MOBILE>
  SPI=1, dst=Localhost, AH type=MD5
```

The table below describes the information elements.

Element	Description
`flags`	Specifies the state and characteristics of the SA. This field can have one of the following values: `ACTIVE`—The SA is active. `INACTIVE`—The SA is inactive. `ESP`—The SA uses an ESP (IP Encapsulating Security Payload, RFC 1827) encryption transform. `AH`—The SA uses an AH (IP Authentication Header, RFC 1826) authentication transform. `MOBILE`—The SA is part of a mobile IP Security scheme. `LOCKED`—The SA is part of a static IP Security scheme.
`SPI`	Specifies the Security Parameters Index (SPI) that the Pipeline uses, along with the destination address, to select the SA.
`dst`	Specifies the destination address associated with the SA. This field can specify one of these values: `local_hostname`—Specifies the name of the local router. `dest_ipaddr` —Specifies the name of the destination host. `Mobile (Unset)`—Specifies an unitialized mobile SA.
`type`	Specifies the type of transform in use: MD5—Indicates the AH-MD5 authentication transform (RFC 1828). SHA1—Indicates the AH-SHA-1 authentication transform (RFC 1852). DES_CBC—Indicates the ESP-DES-CBC (RFC 1829) encryption transform. DES3_CBC—Indicates the ESP-3DES-EDE-CBC (RFC 1851) encryption transform.

Element	Description
`IV size`	Specifies the number of bits in the initialization vector for the ESP-3DES-EDE-CBC encryption transform.
`IV`	Specifies the 3DES key.

IPsecSchemeDump

This command displays one or all IP Security schemes currently residing in the database. You enter the command using this syntax:

```
>IPsecSchemeDump integer
```

If you specify the `integer` argument when entering this command, the Pipeline displays the scheme indicated by the number. If you do not specify an integer, the command displays all schemes in the database.

When you enter the command, this type of information displays:

```
SCHEME 1 at 0x2c3d10:

   flags=1 <ACTIVE>

   dst=Unset(Mobile)

   Incoming AH SA at 0x2c3d50, Outgoing AH SA at 0x2c3e10

   Incoming ESP SA at 0x2c3ed0, Outgoing ESP SA at 0x2c3f90
```

The table below describes the information elements.

Element	Description
`flags`	Specifies the state of the scheme. This field can have one of the following values: `ACTIVE`—The scheme is active. `INACTIVE`—The scheme is inactive.

Element	Description
`dst`	Specifies the destination address associated with the scheme. This field can specify one of these values: `dest_ipaddr` —Specifies the name of the destination host. `Mobile (Unset)`—Specifies an unitialized mobile SA.
`Incoming...` message	Specifies the Security Associations (SAs) that the router receives, as well as the internal router memory address in which the data is stored. ESP indicates that the SA uses an ESP (IP Encapsulating Security Payload, RFC 1827) encryption transform. AH indicates that the SA uses an AH (IP Authentication Header, RFC 1826) authentication transform.
`Outgoing...` message	Specifies the Security Associations (SAs) that the router transmits, as well as the internal router memory address in which the data is stored.

IPsecdblog

This command toggles the diagnostic message status for IP security. You enter the command using this syntax:

```
> IPsecdblog d|s y|n
```

The table below describes the information elements.

Argument	Description
d y	Specifies that the Pipeline write the IP Security diagnostic messages to the diagnostic monitor.
d n	Specifies that the Pipeline does not write IP Security diagnostic messages to the diagnostic monitor.
s y	Specifies that the Pipeline writes IP Security diagnostic messages to Syslog.

Argument	Description
s n	Specifies that the Pipeline does not write IP Security diagnostic messages to Syslog.

IP Security syslog and diagnostic messages

This release includes the following new messages. They can appear in syslog or in the diagnostic monitor.

Message	Description
`IPSEC: Scheme` *num*`: Expected remote` *x.x.x.x*`, got` *y.y.y.y*`.`	The router received a packet with SPIs matching scheme *num*, but the packet came from *x.x.x.x* instead of scheme *num*'s configured tunnel address of *y.y.y.y*.
`IPSEC: Scheme` *num1*`: Expecting SPI` *num2*`, got SPI` *num3*	The router received a packet with both AH and ESP encapsulations, but the two SPIs do not match.
`IPSEC: Scheme` *num1*`: Expecting no encryption, got SPI` *num2*`.`	The router received a packet with SPI *num2*, AH on the outside, and ESP on the inside, but scheme *num1* does not have encryption configured.
`IPSEC: Failed encap on inactive scheme` *num*	The router received a packet for scheme *num* while updating that scheme in the VT100 interface. This message indicates a transient condition.
`IPSEC: Scheme` *num*`: Bogus ICMP Destination Unreachable: Source Route Failed`	The host at the remote end is in an error condition. The router does not send source-routed IPSEC packets, so it should not be receiving this ICMP error message.
`IPSEC: SHA1: Received invalid AH: SPI` *num* *x.x.x.x* `->` *y.y.y.y*	The router received an AH-SHA packet, but the authentication did not match the one configured in the VT100 interface.
`IPSEC: Received unknown SPI` *num*`,` *x.x.x.x* `->` *y.y.y.y*	Received an AH- or ESP-encapsulated packet, but SPI *num* is not configured on the Pipeline.

Message	Description
`IPSEC: Scheme` *num1*`: Expecting no authentication, got SPI` *num2*	Received a packet with SPI *num2*, AH on the outside, and ESP on the inside, but scheme *num1* does not have authentication configured.
`IPSEC: Scheme` *num*`: Received packet encapsulation does not match scheme`	Scheme *num* is configured to use both AH and ESP encapsulation, but the packet used only one type of encapsulation.
`IPSEC: Scheme` *num*`: Bogus ICMP Destination Unreachable: Port Unreachable`	The host at the remote end is in an error condition. The ICMP message refers to bad TCP or UDP port numbers, but the router is sending AH and ESP packets, which do not contain port numbers.
`IPSEC: MD5: Received invalid AH: SPI` *num x.x.x.x -> y.y.y.y*	The router received an AH-MD5 packet, but the authentication did not match the one configured in the VT100 interface.
`IPSEC: Scheme` *num1*`: SPI` *num2* `is not AH,` *x.x.x.x -> y.y.y.y*	The router received an SPI *num2* AH packet, but SPI *num2* is configured for ESP.
`IPSEC: Scheme` *num1*`: SPI` *num2* `is not ESP,` *x.x.x.x -> y.y.y.y*	The router received an SPI *num2* ESP packet, but SPI *num2* is configured for AH.
`IPSEC: Scheme` *num*`: Tunnel endpoint` *x.x.x.x* `is > 64 hops away`	You have misconfigured your scheme to indicate a tunnel endpoint the router cannot find.
`IPSEC: Scheme` *num*`: Tunnel endpoint` *x.x.x.x* `doesn't speak AH`	You have misconfigured your scheme to indicate a tunnel endpoint that cannot handle AH encapsulation.

Message	Description
`IPSEC: Scheme` *num*`: Tunnel endpoint` *x.x.x.x* `doesn't speak ESP`	You have misconfigured your scheme to indicate a tunnel endpoint that cannot handle ESP encapsulation.
`IPSEC: Scheme` *num*`: Tunnel endpoint` *x.x.x.x* `is unreachable`	You have misconfigured your scheme to indicate a tunnel endpoint the router cannot reach.

Firewall changes

This section describes changes to the Ascend firewall.

FWALLversion diagnostic command modified

The FWALLversion diagnostic command now returns a space-delimited list of all firewall versions that the router accepts. If IP security is enabled, the version information is followed by the letter *i*. For example, a current router displays this information:

```
> FWALLversion

FWversion: 1 2 i
```

The router accepts version 1 or version 2 firewalls. IP security can be enabled for both.

Version changes

The firewall version number is now 2, although the Pipeline accepts version 1 firewalls and they work properly.

Language changes

The firewall language saved by SAM (the Secure Access Manager program for Windows) in.fw files has four new keywords. You can use these keywords to control how the Pipeline applies IP Security transforms to incoming and outgoing

packets. The table below describes each keyword.

Keyword	Description
scheme=[*scheme*]	Encapsulate transmitted packets using the information contained in `scheme`. On reception, match the received packet against the parameters in `scheme`, matching only if all parameters match. If no scheme argument appears, the Pipeline uses the scheme remembered using a dynamic trigger rule. Use this keyword without the `scheme` argument only in dynamic rule templates.
auth	Match only received packets that have been successfully authenticated.
crypt	Match only received packets that were successfully decrypted.

Firewall syslog message changes

The one-line message summary lines that display when a firewall logs a packet can now display IP Security-specific information:

- Logged IP packets with a protocol field of 50 display as *esp* instead of *50*.
- Logged IP packets with a protocol field of 51 displays as *ah* instead of *51*.
- Received IP packets that have been successfully authenticated display with the addition of the string *auth* to the end of the logged message.
- Received IP packets that have been successfully decrypted display with the addition of the string *crypt* to the end of the logged message.

Removing down routes to a host

In previous releases, the routes associated with Connection profiles were always advertised and added to the Pipeline unit's routing table, even when the connection was down. This situation posed a problem in some situations, especially for users of nailed-up lines; these users did not want to advertise the routes when the

connection was down. In this release, you can specify that the Pipeline remove routes from the routing table when the associated connection is off line, and not advertise these routes.

Overview

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

However, there are some situations in which this advertisement causes problems. For a nailed-up line, one assumes that the connection is always up. When it is not, the routes to that connection are not necessary until the connection comes up again. The following example illustrates the problem:

Pipeline1 and Pipeline2 are on the same local LAN.

- Pipeline1 has a nailed-up line to a remote site. The remote address has a metric of 4.
- Pipeline2 is a backup connection. It has a remote address with a metric of 7.

Traffic goes through Pipeline1 because of the lower metric. If Pipeline1's nailed-up connection goes down, its route to the remote network is still advertised by default. Therefore, the connection specified by Pipeline2 never comes up.

The Temporary parameter lets you specify that the Pipeline remove the route to an inactive connection's address. The Temporary parameter appears in the IP Options submenu of the Connection profile.

User interface changes

Temporary

Description: Specifies whether the Pipeline stops advertising the route to the address in this Connection profile when the session terminates, and whether the Pipeline removes this route and all routes dynamically learned on this connection from the routing table.

Usage: You can specify one of these settings:

- Yes removes a route from the routing table to a connection when the link is off-line, including all routes dynamically learned on this connection, and discontinues advertising the route.

  The routes are advertised and appear in the routing table only when you re-establish the connection.

- No continues to advertise the route to the connection found in the LAN Adrs and WAN Alias parameters, even when the connection is off-line.

  The route appears in the Pipeline's routing table, along with all other routes dynamically learned on this connection. All routes age normally. The default value is No.

Example: Pipeline1 has a nailed connection with an address of 128.50.69.69. Pipeline1 advertises this route when the connection is up. Pipeline1 also learns through RIP that the remote side is advertising 198.5.248.72. If the connection goes down and Temporary=Yes, the Pipeline removes 128.50.69.69 and 198.5.248.72 from its routing table and no longer advertises them. If the connection goes down and Temporary=No, the Pipeline maintains 128.50.69.69 in the routing table (pointing to the idle interface—wanidle), and allows 198.5.248.72 to age normally.

Dependencies: A frame relay link is a nailed-up connection defined in a Connection profile. A frame relay link can also have a designated backup Connection profile; if the link goes down, the Pipeline uses the backup profile for the connection. To specify the backup profile, you use the Backup parameter. For frame relay links, the effect of the Temporary parameter varies depending upon whether the link has an associated backup profile:

- If a frame relay connection goes down and the frame relay link has a backup Connection profile, the Pipeline ignores a setting of Temporary=Yes.

  The Pipeline does not remove routes from the routing table when the frame relay connection goes down.

- If a frame relay connection goes down and the Frame Relay profile does not have a backup Connection profile, the Pipeline follows a setting of Temporary=Yes.

  The Pipeline removes routes from the routing table when the frame relay connection goes down.

Parameter Location: Connection profile > IP Options > Temporary.

IP routing display changes

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

```
ascend% show ip routes
Destination        Gateway         IF       Flg   Pref
Met     Use     Age
192.168.252.0/30   192.168.252.1   wan10    rGT    60   7
0       7
192.168.252.1/32   192.168.252.1   wan10    rT     60   7
1       7
```

Network summaries for address pools

The following two interfaces, have been added to the routing tables:

- The reject interface (rj0)

  The reject interface has an IP address of 127.0.0.2. Packets routed to this interface are sent back to the source address with an ICMP “host unreachable” message.

- The black-hole interface (bh0)

  The black-hole interface has an IP address of 127.0.0.3. Packets routed to this interface are discarded silently.

Specifying default routes on a per-user basis

You can now specify a default route on a per-user basis in a Connection profile. If an operator is using a particular service, you can cause the Pipeline to send traffic to the router that service uses, even if the router is not the default gateway shown in the system-wide routing table.

Overview

If you specify a default route in a Connection profile, the Pipeline routes IP packets in this way:

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

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

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

Configuration interface changes

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

Client Gateway

Description: Specifies the default route for IP packets coming from the user on this connection.

Usage: Specify the IP address of the next hop router in dotted decimal notation. The default value is 0.0.0.0; if you accept this value, the Pipeline routes packets as specified in the routing table, using the system-wide default route if it cannot find a more specific route.

The Pipeline must have a direct route to the address you specify. The direct route can take place via a profile or an Ethernet connection. If the Pipeline does not have a direct route, it drops the packets on the connection. When you diagnose routing problems with a profile using this feature, an error in a per-user gateway address is not apparent from inspection of the global routing table.

Example: If you specify Client Gateway=10.0.0.3 in a profile, IP packets from the user with destinations through the default route will be routed through the gateway at 10.0.0.3.

Parameter Location: Connection profile: Ethernet > Connections > *any profile* >IP Options

Support for multiple IP routing protocols

This release includes changes to the IP router that provide support for multiple IP routing protocols such as RIP-v1 and RIP-v2.

Route preferences

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

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

  You can modify the default in the Route Preferences submenu of the Ethernet profile.
- A statically configured IP Route or Connection profile has a default Preference=100.

  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 lower number. If the Preference values are equal, the router then compares the Metric field, using the route with the lower Metric.

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

To support route preferences, the following parameters have been added:

Location	Parameter with default value
Ethernet > Connections > *any profile* IP options (Connection profile)	Preference=[]
Ethernet > Static Rtes > *any profile* (IP Route profile)	Preference=[]
Ethernet > Mod Config > Route Ref (Ethernet profile)	Static Preference=100 Rip Preference=100

Preference

Description: Specifies the preference value for a specific statically configured IP route, which may be defined in an IP Route profile or Connection profile. When selecting which routes to put in the routing table, the router first compares the Preference value, selecting the lower number. If the Preference values are equal, then the router compares the Metric field, selecting the route with the lower Metric.

Usage: Press Enter to open a text field. Type a number between 0 and 255. The default value is 100. Zero is the default for connected routes (such as the Ethernet). The value of 255 means “Don't use this route,” which is meaningful only for Connection profiles.

These are the default values for other types of routes:

- Routes learned from ICMP Redirects=30
- Routes learned from RIP=100
- Static routes (configured in an IP Route profile or Connection profile)=100

This set of preference values gives static routes and RIP routes an equal value, with ICMP Redirects taking precedence over both.

Parameter Location: Connection profile: Connections > *profile* > IP Options
IP Route profile: Static Rtes > any profile

Static Preference

Description: Specifies the preference value for statically configured routes created from IP address pools and the Terminal Server IPROUTE ADD command. When selecting which routes to put in the routing table, the router first compares the Preference value, selecting the lower number. If the Preference values are equal, then the router compares the Metric field, selecting the route with the lower Metric.

Usage: Press Enter to open a text field. Then, type a number between 0 and 255. The default value is 100. Zero is the default for connected routes (such as the Ethernet). The value of 255 means “Don't use this route.”

These are the default values for other types of routes:

- Routes learned from ICMP Redirects=30
- Routes learned from RIP=100
- Static routes in an IP Route profile or Connection profile=100

Parameter Location: Ethernet profile: Ethernet > Mod Config > Route Pref

Rip Preference

Description: Specifies the preference value for routes learned from the RIP protocol. When selecting which routes to put in the routing table, the router first compares the Preference value, selecting the lower number. If the Preference values are equal, then the router compares the Metric field, selecting the route with the lower Metric.

Usage: Press Enter to open a text field. Then, type a number between 0 and 255. The default value is 100. Zero is the default for connected routes (such as the Ethernet). The value of 255 means “Don't use this route.”

These are the default values for other types of routes:

- Routes learned from ICMP Redirects=30
- Static routes from IP address pools, RADIUS authentication, and the Terminal Server IPROUTE ADD command=100
- Static routes in an IP Route profile or Connection profile=100

Parameter Location: Ethernet profile: Ethernet > Mod Config > Route Pref

Routing table display changes

The IPROUTE SHOW terminal server command has been modified slightly to include more information relevant to multiple IP routing protocols. To view the IP routing table, invoke the terminal server interface and enter this command at the prompt:

```
iproute show
```

The output looks similar to this:

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

The routes in this table are explained as follows:

- `0.0.0.0/0  10.0.0.100   wan0   SG   1     1     0     20887`
  This is the default route, pointing through the active Connection profile. The IP Route profile for the default route specifies a Preference of 1, so this route is preferred over dynamically learned routes.
- `10.207.76.0/24 10.207.76.1 wanidle0 SG 100 7  0 20887`
  `10.207.76.1/32 10.207.76.1 wanidle0 S  100 7  2 20887`

  These routes are specified in a Connection profile. Note that there are two routes—a direct route to the gateway itself and a route to the larger network.
  `10.207.77.0/24 10.207.76.1 wanidle0 SG 100 8  0  20887`
  This is a static route that points through an inactive gateway.
- `127.0.0.1/32    -       lo0      CP   0     0     0   20887`

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

- `10.0.0.0/24 10.0.0.100 wan0 SG 100 1 21387 20887`
  `10.0.0.100/32 10.0.0.100 wan0 S 100 1 153 20887`

  These routes are created by a Connection profile that is currently active. These are similar to the 10.207.76.0 routes shown above, but these routes live on an active interface.
- `10.1.2.0/24 - ie0 C 0 0 19775 20887`

  This route describes the connection to the Ethernet interface. It is directly connected, with a Preference and Metric of zero.
- `10.1.2.1/32 - lo0 CP 0 0 389 20887`

  This is another loopback route, a host route with our Ethernet address. It is private, so it will not be advertised.
- `255.255.255.255/32 - ie0 CP 0 0 0 20887`

  This is a private route to the broadcast address. This route 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.

Information in the routing table display

The columns in the routing table display the following information:

- Destination

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

  The Gateway column specifies the address of the next hop router that can forward packets to the given destination. Direct routes (without a gateway) no longer show a gateway address in the gateway column.
- IF

  The Interface column shows 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 where all routes point when their WAN connections are down).

- Flg

  The Flg column can contain the following flag values:

  – C=Connected (A directly connected route, for example, the Ethernet.)

  – I=ICMP (ICMP Redirect dynamic route.)

  – N=NetMgt (Pleased in the table via SNMP MIB II.)

  – O=OSPF (A route learned from OSPF.)

  – R (A RIP dynamic route.)

  – S=Static (A locally configured IP Route profile or Connection profile route.)

  – ?=Unknown (A route of unknown error, which indicates an error.)

  – G=Gateway (A gateway is required in order to reach this route.)

  – P=Private (This route will not be advertised via RIP or OSPF.)

  – T=Temporary (This route will be destroyed when its interface goes down.)

  – *=Hidden (A hidden route means that there is a better route in the table, so this route is hidden “behind” the better route. If the better route should go away, then this route may be used.)

  Note that the H (host route) flag has been removed because it was redundant with a /32 netmask in the Destination column. The U (up) flag has also been removed. Physical interfaces are considered “up” once they have been defined in the Ascend Enterprise MIB, so the U flag was contradictory. The D (dynamic route) flag has been replaced by the I (ICMP Redirect) and R (RIP) flags, which are new.

- Pref

  The Preference column contains the preference value of the route. Note that all routes that come from RIP will have a preference value of 100, while the preference value of each individual static route may be set independently.

- Metric

  The Metric column shows the RIP-style metric for the route, with a valid range of 0-16.

- Use

  This is a 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 this route.)

  Note: Unused routes are now indicated by a 0 in the Use column. They were indicated previously by a 1 in the Use column.

- Age

  This is the age of the route in seconds. It is used for troubleshooting, to determine when routes are changing rapidly or flapping.

Routing table and diagnostic changes

Changes have been made to the Pipeline's IP routing stack which improve performance and add additional support for multicast routing. These changes include:

- User interface changes
- Diagnostic mode changes
- Changes to Secure Access Firewall operation

In addition, Pipelines now conform more closely to RFC1812 (Requirements for routers) section 5.

User interface changes

The `iproute show` command output has been modified.

A route has been added from the 127 network to the blackhole interface:

```
127.0.0.0/8    -     bh0     CP     0    0        0       59593
```

Packets routed to the blackhole interface are discarded silently.

Routes pointing to local machines are now labeled local. These include the following routes:

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

with a single w.x.y.z route for each local IP address.

Note that the routes to 224.0.0.1 and 224.0.0.2 are new routes. They represent the multicast addresses for all systems on the local subnet and all routers on the local subnet, respectively, and are never forwarded.

A new route has been added to a virtual interface called mcast. All multicast addresses (except for special addresses such as 224.0.0.1/32 and 224.0.0.2/32) point to the mcast interface:

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

Diagnostic mode changes

The Ippacket diagnostic output has been changed.

Modified diagnostic messages

The wording has been changed in these error messages:

IP: no ip address for this port	IP: received packet on unconfigured interface
IP: options: calling icmp_send(): type = %d code = %d	IP: options: sending icmp to %s, type = %d code = %d
IP: passed pkt length is short	IP: received frame too small to hold any IP header
IP: short IP header	IP: received packet with header size < 20 bytes
IP: version check failed	IP: received unknown IP version %d
IP: bootp packet	IP: received BOOTP packet
IP: NAT packet	IP: received NAT packet
IP: checksum failed	IP: received bad checksum
IP: no memory	IP: no memory, dropping packet

New diagnostic messages

The following messages have been added:

- IP: received packet too small to hold its IP header
- IP: received truncated IP packet
- IP: received 0 ttl

Deleted diagnostic messages

The following messages have been deleted:

- IP: passed pkt length is short
- IP: (pkt <MIN_ETHER_LEN) length check failed
- IP: (pkt >MAX_ETHER_LEN) length check failed
- IP: (pkt <=MAX_ETHER_LEN) length check failed
- IP: (pkt <=MAX_ETHER_LEN) is padded
- IP: short length check failed
- IP: IF wants gateway %s, but no route
- IP: route to gateway %s isn't direct
- IP: (next hop to it is %s)
- IP: no route to %s.
- IP: not forwarding
- IP: bad incoming ttl of zero!!!
- IP: ttl expired
- Bad checksum pkt at 0x%p
- IP: parse: not bcast
- IP: parse: source & dest if different
- IP: NAT Session not active
- IP: reassembly error
- IP: not joined
- IP: unused Pool address.

Changes to Secure Access Firewall operation

A minor change has been made in the way that a Secure Access Firewall deals with directed broadcasts. A directed broadcast, received as a unicast, will not be

delivered locally if the firewall on the outbound interface would block that packet. Thus, if the firewall on the outbound interface is set up to block a packet, no one will receive it, including the Pipeline. Previously the packet would be routed by the Pipeline to the outbound interface, where it would be dropped by the firewall.

Interface-based routing

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

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

This feature allows the user to configure each link as numbered (interface-based) or unnumbered (system-based). If no interfaces are specified as numbered, then the unit will operate exactly as it has previously. Interface numbering is accomplished via the Connection profile.

System behavior with a numbered interface

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

- IP packets generated in the Pipeline and sent to the remote address will have an IP source address corresponding to the numbered interface, not to the default (Ethernet) address of the Pipeline.

- During authentication of a call placed from a Pipeline using a numbered interface, the Pipeline will report the address of the interface as its IP address.
- The Pipeline will add as host routes to its routing table, all numbered interfaces listed in Connection Profiles.
- The Pipeline will accept IP packets whose destination is a numbered interface listed in a Connection profile, considering them to be destined for the Pipeline itself. (The packet may actually arrive over any interface, and the numbered interface corresponding to the packet's destination address need not be in the active state.)

Interface IP Address (IF Adrs) parameter

Configuration of a numbered link takes place in the Connection profile, under the IP Options submenu. A new parameter, IF Adrs, specifies the IP address of the interface. If the field is left at its default value (0.0.0.0/0), then the interface will be treated as unnumbered.

The profile below shows a typical profile for an unnumbered interface. The new IF Adrs field is not used for an unnumbered interface.

```
90-103
  Ip options...
  LAN Adrs=192.168.6.29/24
  WAN Alias=0.0.0.0/0
  IF Adrs=0.0.0.0/0
  Metric=0
  reference=2
  Private=No
  IP=Off
  Pool=0
```

The profile below shows settings for a numbered interface. The WAN Alias parameters has been filled in with the address of the remote end of the link, and the new IF Adrs parameter contains the number of the interface at the near end of the link.

```
90-103
 Ip options...
 LAN Adrs=192.168.6.29/24
 WAN Alias=192.1.1.17
 IF Adrs=192.1.1.8/30
```

```
Metric=0
Preference=2
Private=No
RIP=Off
Pool=0
```

IF Adrs

Description: Specifies the IP address of the interface at the near end of a link.

Usage: Press Enter to open a text field. Then, type the IP address of the numbered interface.

An IP address consists of four numbers between 0 and 255, separated by periods. If a netmask is in use on the network, you must specify it. Separate the netmask from the IP address with a slash. The default is 0.0.0.0/0.

Press Enter again to close the text field.

Example: `200.207.23.7/24`

Dependencies: The IF Adrs parameter does not apply if the Pipeline does not support IP (Route IP=No).

Parameter Location: Connection Profile: Ethernet > Connections > IP options

See Also: WAN Alias, Route IP

Specifying the remote interface address

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

If both the system and interface addresses are known

If interface-based routing is being added to a system which has already been set up using system-based routing, the easiest way to specify the remote interface address is by using the WAN Alias parameter in the Connection profile. WAN Alias is used to identify the remote end of the link. If a WAN Alias is set, the following will take place:

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

If you want to create static routes to hosts at the remote end, you can use the WAN Alias address as the “next hop” (gateway) field. (The Lan Adrs address will also work, as would be used for system-based routing.)

If only the interface address is known

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

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

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

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

If the remote interface address is not specified

If interface-based routing is in use and the local interface is numbered, the remote address will usually be known (in practice, the subnet must be agreed upon by administrators of both sites.) It is possible, but not recommended, to number the local interface, omitting the interface address of the remote site and using only its

system or LAN address. In that case, do not use the (supposedly unknown) remote interface address in any static routes.

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

Multicast forwarding and IGMP functionality

The Pipeline now supports Internet Group Membership Protocol (IGMP) version 1 and version 2, along with configuration options that enable the Pipeline to forward multicast traffic. The Pipeline communicates with a multicast router on its Ethernet interface and forwards multicast traffic to dial-in multicast clients connected via Pipelines. The Pipeline transparently passes multicast traffic between the multicast router and its clients.

Configuring the Pipeline for multicast forwarding

To configure the Pipeline for multicast forwarding:

1 Select Ethernet > Mod Config > Multicast.

2 Set Forwarding to Yes.

3 Specify the Multicast Profile used to connect to the multicast router.

4 Reboot the Pipeline for the changes to take effect.

Parameter reference

Two new parameters have been added to support the multicast feature:

- Ethernet > Mod Config > Multicast Forwarding
- Ethernet > Mod Config > Multicast Profile

This section describes these new parameters.

Multicast Forwarding

Description: Enables multicast forwarding in the IP-only version of the Pipeline. By default, it is set to No.

Usage: Press Enter to toggle between Yes and No.

- Yes enables multicast forwarding.

  When set to Yes, the Pipeline appears to a Multicast router as a multicast client, which receives Internet Group Membership Protocol (IGMP) queries from the router and responds to them using IGMP. To dial-in clients, it appears as a multicast router, which sends IGMP queries and forwards multicast traffic.
- No, the default, turns off multicast forwarding.

Parameter Location: Ethernet > Mod Config > Ethernet Profile

Dependencies: Available in the IP-only release for the Pipeline.

See Also: Multicast Profile

Multicast Profile

Description: Specifies the name of a Connection Profile for a WAN link to a multicast router in the IP-only version of the Pipeline. If no profile name is specified and Multicast Forwarding is turned on, the Pipeline assumes that its Ethernet is the Multicast interface.

The specified Connection Profile must be resident.

Usage: Press Enter to open a text field. Then, type the name of the Connection Profile to the multicast interface. If no name is specified, the Pipeline assumes the presence of a multicast router on its Ethernet interface. Press Enter again to close the text field.

Parameter Location: Ethernet > Mod Config

Dependencies: Available only in the IP-only release for the Pipeline. It is not available if Multicast Forwarding is set to No.

See Also: Multicast Forwarding

Terminal server commands for multicast forwarding and IGMP

The Show commands described in this section have been added to the terminal server command line to support multicast functionality. To use them, first invoke the terminal server interface (System > Sys Diag> Term Serv).

To display all the active multicast group addresses and the clients(interfaces) registered for that group, type:

```
ascend% show igmp groups
```

The output is similar to this:

```
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
```

Where:

- Hash is an index to a hash table (displayed for diagnostics purposes only).
- Group address is the IP multicast address used in this packet.
- 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 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. When this field contains periods, it means that this membership never expires.
- Counts shows 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 is displayed for diagnostics purposes).

To list all IGMP multicast clients, enter:

```
ascend% show igmp clients
```

The output is similar to this:

```
IGMP Clients

 Client      Version  RecvCount  CLU      ALU
  0(Mbone)      1        0         0       0
  2             1       39        68       67
  1             1    33310        65       65
```

Where:

- Client indicates the interface ID on which the client 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.
- Version is the version of IGMP being used.
- RecvCount is the number of IGMP messages received on that interface.
- CLU (current line utilization) and ALU (average line utilization) show the percentage of bandwidth utilized across this interface. If bandwidth utilization is high, some IGMP packet types will not be forwarded.

To display IGMP activity statistics, type:

```
ascend% show igmp stats
```

The output shows the number of IGMP packet types sent and received, in the format below:

```
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.
```

IP Address Management

4

Overview

The following new features might affect the way you assign IP address for users on your LAN:

Network Address Translation (NAT) for a LAN

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

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

NAT can be implemented to use a single address or multiple addresses. Using multiple IP addresses requires a remote MAX configured as a DHCP server and a version 2 Pipeline 75 BRI unit.

Single-address NAT and port routing

A Pipeline 50, 75, or 130 can perform single-address NAT in these ways:

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

Note: If you have a version 2 Pipeline 75 BRI, you can use single-address NAT by setting the Ethernet > NAT > Lan parameter to Single IP Addr. For all other units, single-address NAT is the default and the Lan parameter is hidden.

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

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

  Any number of hosts on the local network can make any number of simultaneous connections to hosts on the remote network, subject only to the memory limitations of the Pipeline. The translations between the local network and the Internet or remote network are dynamic and do not need to be preconfigured, as do incoming connections.

- For incoming calls, the Pipeline can perform NAT for multiple hosts on the local network using its own IP address. The Pipeline routes incoming packets for up to 10 different TCP or UDP ports to specific servers on the local network. Translations between the local network and the Internet or remote network are static and need to be preconfigured. You need to define a list of local servers and the UDP and TCP ports each would handle. You can also define a local default server that handles UDP and TCP ports not listed.

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

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

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

About the NAT Profile

In the NAT profile you specify the name of a Connection profile for which the Pipeline performs network address translation on incoming and outgoing calls. You also specify other parameters needed by NAT.

Port routing and firewalls

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

Configuring single-address NAT and port routing

You can configure a Pipeline to perform NAT for remote hosts wishing to access services on the local network in one of these ways:

- Route all incoming packets from a remote network to a single server on the local network.
- 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.

The following sections explain how.

Routing all incoming packets to a single server

To configure the Pipeline to perform NAT and route all incoming packets from a remote network to a single server on the local network:

1 Open the Ethernet > NAT menu.

2 Set the Routing parameter to Yes.

3 Set the Profile parameter to the name of an existing Connection Profile.

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

4 If you're configuring a version 2 Pipeline 75 BRI, set the Lan parameter to Single IP Addr.

5 If you previously configured the Pipeline to route incoming packets for specific TCP or UDP ports (as described in "Routing incoming packets for specific ports" on page 4-5):

   – Open each Ethernet > NAT > Static Mapping > Static Mapping *nn* menu (where *nn* is a number between 01 and 10).

   – Set the Valid parameter in each menu to No.

6 Set the Def Server parameter to the IP address of the server on the local network to receive all incoming packets from the remote network.

7 Press the Esc key to exit the menu.

8 Save the changes when prompted.

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

Routing incoming packets for specific ports

To 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:

1 Open the Ethernet > NAT menu.

2 Set the Routing parameter to Yes.

3 Set the Profile parameter to the name of an existing Connection Profile.

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

4 If you're configuring a version 2 Pipeline 75 BRI, set the Lan parameter to Single IP Addr.

5 Open the Ethernet > NAT > Static Mapping menu.

6 Open a Static Mapping *nn* menu, where *nn* is a number between 01 and 10.

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

7 Set the Valid parameter to Yes.

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

8. Set the Dst Port# parameter to the number of a TCP or UDP port. See "Well-known ports" on page 4-7 for information on obtaining port numbers.

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

9. Set the Protocol parameter to TCP or UDP.

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

10. Set the Local Port# to the port on the local server to route packets to.

11. Set the Local Adrs parameter to the address of the local server to route packets to.

12. Press the Esc key to exit the menu.

13. Save the changes when prompted.

14. Repeat steps 6 through 13 for any additional ports whose packets you want to route to a specific server and port on the local network.

15. Open the Ethernet > NAT menu.

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

17. Press the Esc key to exit the menu.

18. Save the changes when prompted.

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

Disabling routing for specific ports

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

1. Open the Ethernet > NAT > Static Mapping menu.

2. Open a Static Mapping *nn* menu, where *nn* is a number between 01 and 10.

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

3 Set the Valid parameter to No.

   This disables routing for the port specified by the Dst Port# and Protocol parameters in this menu.

4 Press the Esc key to exit the menu.

5 Save the changes when prompted.

6 Repeat steps 2 through 5 to disable routing for any additional ports.

7 Press the Esc key to exit the menu.

8 Save the changes when prompted.

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

Well-known ports

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

You can obtain current lists of Well Known Ports and Registered Ports (ports in the range 1024-4915 that have been registered with the IANA) via FTP from

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

Multiple-address NAT

For the version 2 Pipeline 75 BRI, multiple-address NAT can be performed when translating addresses for more than one host on the local network. To do this, the Pipeline borrows an official IP address for each host from a Dynamic Host Configuration Protocol (DHCP) server on the remote network.

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

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

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

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

When the first client on the LAN requests access to the remote network, the Pipeline gets this address through PPP negotiation. When subsequent clients request access to the remote network, the Pipeline asks for an IP address from the MAX using a DHCP request packet. The MAX then sends an address to the Pipeline from its IP address pool. The Pipeline uses the dynamic addresses it receives from the MAX to translate IP addresses on behalf of local clients.

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

IP addresses are typically offered by the MAX only for a limited duration, but the Pipeline automatically renews the lease on these addresses. If the connection to the remote server is dropped, all leased addresses are considered revoked. Therefore, TCP connections will not persist across calls.

The Pipeline itself does not have an address on the remote network. This means that the Pipeline can only be accessed from the local network, not from the WAN.

In some installations, the MAX will be handling both NAT DHCP requests and

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

Configuring multiple-address NAT

Configuring Multiple-Address NAT requires that you configure both the Pipeline and the MAX it connects to.

To configure NAT on the Pipeline:

1 Open the menu Ethernet > NAT.

2 Set Profile to the name of the Connection Profile use to connect to the MAX (the DHCP server) as described below.

3 Set Routing to Yes, and Lan to Multiple.

This avoids the possibility that a default route from the ISP will overwrite the NAT route.

4 Exit the Mod Config menu and save your changes.

Configuring NAT on the MAX

Note: The information in this section is for your convenience only, and is meant to help you work with the administrator of the MAX DHCP server. For more detail, or more up-to-date information, please refer to the MAX documentation about configuring the MAX as a DHCP server.

Where to Configure NAT on the MAX DHCP server:

- If you are using settings in the Answer profile to build the connection using the Use Answer as Defaults parameter (for RADIUS) or if you are using Names/Passwords profiles, configure NAT In the Answer profile.
- If you are using Connection profiles for users or if you are using the Template Connection # parameter (for Names/Passwords profiles), configure NAT in a Connection profile.
- If you are using RADIUS, configure NAT in a RADIUS profile.

To configure NAT:

1 From the main edit menu select:

– Ethernet > Answer > DHCP options *or*

– Ethernet > Connections > *NAT Connection Profile* > DHCP options

2 Set Reply Enabled=Yes.

3 Set Pool Number to the IP address pool to use for allocating IP addresses to NAT clients. Setting Pool Number to 0 indicates that any pool can be used.

4 Set MAX Leases to the number of addresses to be given to the Pipeline.

5 If you use RADIUS to authenticate users and you do not authenticate users that request DHCP, set Use Answer as Defaults to Yes in the Answer profile. Otherwise, the MAX will not act as a DHCP server for these clients.

NAT menus

To access the NAT settings, open the Ethernet > NAT > NAT menu, which contains these parameters.

```
20-A00 NAT
   Routing=Yes
   Profile=max4
   FR address=0.0.0.0
   Lan=Single IP addr
   Static Mappings...
   Def Server=181.81.8.1
   Reuse last addr=No
   Reuse adr timeout=0
```

The Static Mappings menu includes 10 Static Mapping submenus, which are numbered from 01 to 10. Each submenus contains parameters for controlling the routing of packets from a remote network to a specific TCP or UDP port:

```
NAT
  Static Mapping 01
  ...
  Static Mapping 10
```

Each Static Mapping menu contains the following parameters:

```
20-A00 NAT
   Static Mapping 01
      Valid=Yes
```

```
Dst Port#=21
Protocol=TCP
Loc Port#=21
Loc Adrs=181.100.100.102
```

NAT for Frame Relay

The single-IP address implementation of network address translation (NAT) has been extended to work with Frame Relay. A Pipeline 50, 75, or 130 using Frame Relay encapsulation can now use a single IP address to translate local addresses into a single, official address for networking over the wide area network and accessing the Internet.

Defining the Frame Relay address in the NAT menu

When using Frame Relay, the Ethernet > NAT menu is changed as follows:

```
20-A00 NAT
   Routing=Yes
   Profile=max4
   FR address=
   Static Mapping...
   Def Server=181.81.8.1
```

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

Parameter reference

Def Server

Description: When the Pipeline is configured to perform network address translation (NAT), it can route packets from a remote network for up to 10 different TCP or UDP ports to specific servers and ports on the local network. This parameter specifies a local server to which the Pipeline routes any incoming packets that are *not* routed to a specific server and port.

Note: If you change the value of this parameter, the change does not take effect until the next time a connection is made to the remote network specified in the NAT Profile. To make the change immediately, you must terminate the connection to the remote network and then reopen it.

Usage: Press Enter to open a text field and then type the IP address.

The address consists of four numbers between 0 and 255, separated by periods. Enter 0.0.0.0 to disable routing of packets to a default server.

The default value is 0.0.0.0.

Press Enter again to close the text field, press Esc to exit the menu, and then confirm the change when prompted.

Note: The change does not take effect until the next time the link is brought up. To make the change immediately, bring the link down and back up.

Dependencies: Keep this additional information in mind:

- For routing of packets from a remote network to occur, the Routing parameter in the NAT menu must be set to Yes and the Lan parameter in the NAT menu must be set to Single IP Addr. Parameters in Static Mapping *nn* menus (where *nn* is a number between 01 and 10) control whether the Pipeline routes packets from a remote network for up to 10 different TCP or UDP ports to specific servers and ports on the local network.
  - The Dst Port# and Loc Port# parameters must be set to values other than 0.
  - The address can’t be 0.
- If your local network has only one server that handles all incoming packets, you can specify the server by
  - setting this parameter to the address of the server.
  - setting the Valid parameter in each of the Static Mapping *nn* menus to No, which disables routing of incoming packets by their destination ports.
- If the Routing parameter in the NAT menu is set to No or the Lan parameter in the NAT menu is set to Multi IP Addr, this parameter is N/A.

Location: Ethernet > NAT

See Also: Dst Port#, Loc Adrs, Loc Port#, Lan, Routing, Protocol, Valid

---

Dst Port#

Description: This parameter specifies a TCP or UDP port on the Pipeline to which the remote network sends packets. The Pipeline can route packets for this port to a specific server and port on the local network. This routing, which occurs only in conjunction with network address translation (NAT), is controlled by the parameters in the same Static Mapping *nn* menu (where *nn* is a number between 01 and 10).

Note: If you change the value of this parameter or of any of the other parameters in a Static Mapping *nn* menu, the change does not take effect until the next time a connection is made to the remote network specified in the NAT Profile. To make the change immediately, you must terminate the connection to the remote network and then reopen it.

Usage: Press Enter to open a text field and then type the port number.

Enter a port number between 1 and 65535.

Press Enter again to close the text field, press Esc to exit the menu, and then confirm the change when prompted.

Note: The change does not take effect until the next time the link is brought up. To make the change immediately, bring the link down and back up.

Dependencies: Keep this additional information in mind:

- For routing of incoming packets for a particular port to occur, the Routing parameter in the NAT menu must be set to Yes, the Lan parameter in the NAT menu must be set to Single IP Addr, the Valid parameter in the same Static Mapping *nn* menu must be set to Yes, and other parameters in the same Static Mapping *nn* menu must be set to non-null values:
  - The Loc Port# parameter must be set to a value other than 0.
  - The Loc Adrs parameter must be set to an address other than 0.0.0.0.

  If you enter 0 as the value of this parameter, you receive the message Invalid Input: Zero input is not Valid.

- The Protocol parameter in the same Static Mapping *nn* menu determines whether the port you specify is a TCP or UDP port.
- If the Routing parameter in the NAT menu is set to No or the Lan parameter in the NAT menu is set to Multi IP Addr, this parameter is N/A.

Location: Ethernet > NAT > Static Mapping > Static Mapping *nn* (where *nn* is a number between 01 and 10)

See Also: Def Server, Loc Adrs, Loc Port#, Lan, Routing, Protocol, Valid

---

FR address

Description: The single IP address used when translating local addresses into a single, official IP address for networking over the wide area network and accessing the Internet.

Usage: Press Enter to open a text field and then type the official IP address.

The address consists of four numbers between 0 and 255, separated by periods. You must enter a valid IP address for the feature to work.

Dependencies: Keep this additional information in mind:

- In your connection profile, you must set Encaps=FR. For connection that do not use Frame Relay.
- The Routing parameter in the NAT menu must be set to Yes.

Parameter Location: Ethernet > NAT

See Also: Encaps, Routing, Profile, Static Mappings, and Def Server.

Lan

Description: Available only on the version 2 Pipeline 75 BRI. The default is Single IP Addr.

Usage: Press Enter to toggle between Single IP Addr and Multiple IP Addr. In Single IP Addr, the only host that is visible to the remote network is the Pipeline. In Multiple IP Addr, all the hosts that receive an IP address from the MAX (DHCP server) are visible to the remote network. In Single IP Addr, the Pipeline

receives an IP address on outgoing calls during PPP negotiations and uses its own IP address on incoming calls. Multiple IP Addr does not support incoming calls.

Dependencies: For an incoming connection to initiate NAT, the value of the Lan parameter must be Single IP Address.

Parameter Location: Ethernet > NAT

Loc Adrs

Description: When the Pipeline is configured to perform network address translation (NAT) and to route packets for a particular TCP or UDP port it receives from a remote network to a specific server and port on the local network, this parameter specifies the server to which to route the packets.

Note: If you change the value of this parameter or of any of the other parameters in a Static Mapping *nn* menu, the change does not take effect until the next time a connection is made to the remote network specified in the NAT Profile. To make the change immediately, you must terminate the connection to the remote network and then reopen it.

Usage: Press Enter to open a text field and then type the IP address.

The address consists of four numbers between 0 and 255, separated by periods. Enter 0.0.0.0 to disable routing of packets.

The default value is 0.0.0.0.

Press Enter again to close the text field, press Esc to exit the menu, and then confirm the change when prompted.

Note: The change does not take effect until the next time the link is brought up. To make the change immediately, bring the link down and back up.

Dependencies: Keep this additional information in mind:

- For routing of incoming packets for a particular port to occur, the Routing parameter in the NAT menu must be set to Yes, the Lan parameter in the NAT menu must be set to Single IP Addr, the Valid parameter in the same Static Mapping *nn* menu must be set to Yes, and other parameters in the same Static Mapping *nn* menu must be set to non-null values:

– Dst Port# and Loc Port# parameters must be set to values other than 0.

If you enter 0 as the value of this parameter, you receive the message Invalid Input: Zero input is not Valid.

- If the Routing parameter in the NAT menu is set to No or the Lan parameter in the NAT menu is set to Multi IP Addr, this parameter is N/A.

Location: Ethernet > NAT > Static Mapping > Static Mapping *nn* (where *nn* is a number between 01 and 10)

See Also: Def Server, Dst Port#, Loc Port#, Lan, Routing, Protocol, Valid

---

Loc Port#

Description: When the Pipeline is configured to perform network address translation (NAT) and to route packets for a particular TCP or UDP port it receives from a remote network to a specific server and port on the local network, this parameter specifies the port on the local server to which to route the packets. This port does not have to be the same as the port on the Pipeline to which the packets were originally sent.

Note: If you change the value of this parameter or of any of the other parameters in a Static Mapping *nn* menu, the change does not take effect until the next time a connection is made to the remote network specified in the NAT Profile. To make the change immediately, you must terminate the connection to the remote network and then reopen it.

Usage: Press Enter to open a text field and then type the port number.

Enter a port number between 1 and 65535, or enter 0 to disable routing of packets. 0 is the default.

Press Enter again to close the text field, press Esc to exit the menu, and then confirm the change when prompted.

Note: The change does not take effect until the next time the link is brought up. To make the change immediately, bring the link down and back up.

Dependencies: Keep this additional information in mind:

- For routing of incoming packets for a particular port to occur, the Routing parameter in the NAT menu must be set to Yes, the Lan parameter in the NAT menu must be set to Single IP Addr, the Valid parameter in the same Static Mapping *nn* menu must be set to Yes, and other parameters in the same Static Mapping *nn* menu must be set to non-null values:
  - The Dst Port# parameter must be set to a value other than 0.
  - The Loc Adrs parameter must be set to an address other than 0.0.0.0.

  If you enter 0.0.0.0 as the value of this parameter, you receive the message Invalid Input: Zero input is not Valid.
- The Protocol parameter in the same Static Mapping *nn* menu determines whether the port you specify is a TCP or UDP port.
- If the Routing parameter in the NAT menu is set to No or the Lan parameter in the NAT menu is set to Multi IP Addr, this parameter is N/A.
- You cannot specify the same server and port in more than one Static Mapping *nn* menu.

Location: Ethernet > NAT > Static Mapping > Static Mapping *nn* (where *nn* is a number between 01 and 10)

See Also: Def Server, Dst Port#, Loc Adrs, Lan, Routing, Protocol, Valid

---

Profile

Description: This parameter specifies the name of a Connection Profile used to connect a remote network to the Pipeline. If the Pipeline is configured to perform network address translation (NAT), the Pipeline automatically performs NAT whenever a connection is made with this profile. The profile can be configured for incoming connections, outgoing connections, or both. If the profile is used for an outgoing connection, the remote server must be configured provide valid IP addresses for NAT, either through PPP negotiation for a single address or DHCP for the multiple addresses needed for NAT for LAN.

Usage: Press Enter to open a text field and then enter the name of a Connection Profile.

Press Enter again to close the text field, press Esc to exit the menu, and then confirm the change when prompted.

> **Note:** The change does not take effect until the next time the link is brought up. To make the change immediately, bring the link down and back up.

Dependencies: Keep this additional information in mind.

- If the Routing parameter in the NAT menu is set to No, this parameter is N/A.
- If you specify a Connection Profile that does not exist, the Pipeline does not perform NAT.

Location: Ethernet > NAT

See Also: Routing

Protocol

Description: This parameter specifies whether the Dst Port# and Loc Port# parameters in the same Static Mapping *nn* menu (where *nn* is a number between 01 and 10) specify TCP or UDP ports.

Note: If you change the value of this parameter or of any of the other parameters in a Static Mapping *nn* menu, the change does not take effect until the next time a connection is made to the remote network specified in the NAT Profile. To make the change immediately, you must terminate the connection to the remote network and then reopen it.

Usage: Press Enter to toggle between the choices, press Esc to exit the menu, and then confirm the change when prompted.

- TCP specifies that the Dst Port# and Loc Port# parameters in the same Static Mapping *nn* menu are TCP port numbers.

  TCP is the default.
- UDP specifies that the Dst Port# and Loc Port# parameters in the same Static Mapping *nn* menu are UDP port numbers.

  Note: The change does not take effect until the next time the link is brought up. To make the change immediately, bring the link down and back up.

Dependencies: If the Routing parameter in the NAT menu is set to No or the Lan parameter in the NAT menu is set to Multi IP Addr, this parameter is N/A.

Location: Ethernet > NAT > Static Mapping > Static Mapping *nn* (where *nn* is a number between 01 and 10)

See Also: Dst Port#, Loc Port#

Reuse last addr

Description: Specifies that the last IP address given by the DHCP server should be reused in subsequent DHCP negotiations (for the duration specified in the Reuse addr timeout parameter). Set this parameter when you need to use the same IP address for TCP applications that do not time out, such as Telnet. When the WAN session is idle for a long enough period of time, it will timeout, and the next time the WAN session is established, a new IP address will be assigned by the DHCP server. This creates a problem for users of applications that don't time out, since they expect to be using the same IP address.

Usage: The possible values are Yes or No. The default is No.

Dependencies: This setting is N/A if NAT routing is disabled, or when using Multiple-address NAT. Additionally, it does not apply if the value of Reuse addr timeout has been reached.

If the original IP address given during DHCP negotiation is lost, and cannot be reused, applications requiring the same IP address will need to be reset.

Parameter Location: Ethernet > NAT

See Also: Reuse addr timeout

Reuse addr time-out

Description: Specifies the number of minutes to lease the IP address obtained during DHCP negotiation when Reuse last addr is set to Yes. During the period of time set in this parameter, even if the WAN session is idle and times out, the same

address will be associated with the WAN connection each time it is re-established.

Usage: Set the value to a number of minutes, from 0 to 1440.

- When set to 0, the timer is disabled.
  The default is 0.
- The maximum setting is 1440, which is 24 hours.

Dependencies: This parameter is N/A if NAT routing is disabled, or if NAT is enabled, but Multiple-address NAT is being used.

Parameter Location: Ethernet > NAT

See Also: Reuse last addr

Routing

Description: This parameter enables or disables network address translation (NAT).

NAT is a service provided to one or more hosts on the local network that do not have official IP addresses for a remote network. It works as follows:

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

NAT can be configured to work in one of two ways:

- The Pipeline makes a connection to a remote network, gets an official IP address for the remote network through PPP negotiation, and uses the official address for all address translations.
- The Pipeline makes a connection to a remote network, borrows multiple official IP addresses from a DHCP server on the remote network, and uses each address for translating the packets to and from a specific host on the local network.

When NAT is disabled, the Pipeline releases any IP addresses it has borrowed from the remote network, translation stops, and packets flow between LAN and WAN as they normally would.

Usage: Press Enter to toggle between Yes and No, press Esc to exit the menu, and then confirm the change when prompted.

- Yes enables NAT.
- No disables NAT.

  No is the default.

  Note: The change does not take effect until the next time the link is brought up. To make the change immediately, bring the link down and back up.

Dependencies: Keep this additional information in mind:

- To use NAT, IP routing must be enabled on the Pipeline.
- The IP addresses of hosts on the local network that use NAT and the Pipeline must be on the same subnet. These addresses are only used for local communication between the host and the Pipeline over the Ethernet.
- You should restrict IP addresses used on the local LAN so that hosts on the network connecting to the Pipeline have each octet of their IP addresses greater than 99 (this only applies to FTP sessions). For example, 192.168.121.101 is a recommended address, but 192.168.121.99 is not.
- When the Pipeline connects to a remote network, the MAX or other remote device must be configured to assign dynamic IP addresses through PPP negotiations (when the Pipeline uses a single IP address for NAT) or through DHCP (when the Pipeline requires multiple IP addresses when performing NAT for a LAN).
- Once a connection is terminated, there is no guarantee that the same IP address will be used for subsequent connections. You can set the Idle timer (in the Sessions options submenu of the Connection Profile) to 0 to prevent the Pipeline from terminating an idle connection. But note that the MAX or other device on the remote network may have the Idle timer configured to a lower value, which overrides any settings you have set.
- Once NAT has been configured and the Pipeline is translating addresses from clients on the local LAN, the Pipeline can only be accessed from the local LAN or through the serial port; it cannot be directly accessed from the WAN side.

- Note that the Pipeline itself can be a NAT client. That is, the Pipeline can translate an address for itself as long as it is not translating addresses for other clients on the local LAN.
- Make sure to set Ignore Def Rt=Yes. When NAT is active, it routes using its own default route. Configuring the Pipeline to ignore default routes avoids the possibility that a default route from the ISP will overwrite the NAT route.

Location: Ethernet > NAT

See Also: Def Server, Dst Port#, Loc Adrs, Loc Port#, Lan, Routing, Protocol

---

Valid

Description: This parameter enables or disables the routing of incoming packets for a particular TCP or UDP port to a specific server and port on the local network. This routing, which occurs only in conjunction with network address translation (NAT), is controlled by the parameters in the same Static Mapping *nn* menu (where *nn* is a number between 01 and 10).

Note: If you change the value of this parameter or of any of the other parameters in a Static Mapping *nn* menu, the change does not take effect until the next time a connection is made to the remote network specified in the NAT Profile. To make the change immediately, you must terminate the connection to the remote network and then reopen it.

Usage: Press Enter to toggle between Yes and No, press Esc to exit the menu, and then confirm the change when prompted.

- Yes enables the routing of incoming packets specified by the other parameters in the same Static Mapping *nn* menu.
- No disables the routing of incoming packets specified by the other parameters in the same Static Mapping *nn* menu.

  No is the default.

  Note: The change does not take effect until the next time the link is brought up. To make the change immediately, bring the link down and back up.

Dependencies: For routing of incoming packets for a particular port to occur, the Routing parameter in the NAT menu must be set to Yes, the Lan parameter in

the NAT menu must be set to Single IP Addr, and other parameters in the same Static Mapping *nn* menu must be set to non-null values:

- The Dst Port# and Loc Port# parameters must be set to values other than 0.
- The Loc Adrs parameter must be set to an address other than 0.0.0.0.

Location: Ethernet > NAT > Static Mapping > Static Mapping *nn* (where *nn* is a number between 01 and 10)

See Also: Def Server, Dst Port#, Loc Adrs, Loc Port#, Lan, Routing, Protocol

BOOTP Relay

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

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

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

Using BOOTP relay

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

1 Obtain the IP address of up to two BOOTP server(s) to be used.

2 Open the Ethernet > Mod Config:

```
20-A00 Mod Config
 BOOTP Relay...
```

```
>BOOTP Relay Enable=No
 Server=0.0.0.0
 Server=0.0.0.0
```

3. Select BOOTP Relay Enable and set it to Yes.
4. Select Server and press Enter to open a text box. In the text box, enter the IP address of the BOOTP server. Press Enter to close the text box.
5. If there is another BOOTP server available, select the second menu item named Server and enter its IP address.

   You are not required to specify a second BOOTP server.

   Note: If you specify two BOOTP servers, the Pipeline that relays the BOOTP request determines when each server is used. The order of the BOOTP servers in the BOOTP Relay menu does not necessarily determine which server is tried first.

Parameter reference

BOOTP Relay Enable

Description: Controls whether Bootstrap Protocol (BOOTP) requests are relayed to other networks.

Usage: Press Enter to cycle through the choices.

- Yes specifies that BOOTP requests are relayed.
- No specifies that BOOTP requests are not relayed.

  No is the default.

Dependencies: You must use the Server parameter to specify the address of at least one BOOTP server. The BOOTP Relay menu also includes a second Server parameter for specifying a second BOOTP server. If you specify two BOOTP servers, the Pipeline that relays the BOOTP request determines when each server is used. The order of the BOOTP servers in the BOOTP Relay menu does not necessarily determine which server is tried first.

Parameter Location: Mod Config, BOOTP Relay

See Also: Server, DHCP Spoofing

Server

Description: Specifies a Bootstrap Protocol (BOOTP) server for handling BOOTP requests. If a server is on the same local-area network as the Pipeline, BOOTP requests from other networks are relayed to the server. If a server is on another network, BOOTP requests from clients on the same local-area network as the Pipeline are relayed to the remote server.

Note: This parameter appears twice. Each copy can be used to specify a different BOOTP server.

Usage: Press Enter to open a text field and then type the IP address of the BOOTP server. When you're done, press Enter to close the text field.

Dependencies: If you specify two BOOTP servers, the Pipeline that relays the BOOTP request determines when each server is used. The order of the BOOTP servers in the BOOTP Relay menu does not necessarily determine which server is tried first.

Parameter Location: Mod Config, BOOTP Relay

See Also: BOOTP Relay Enable

DHCP services enhanced

A Pipeline 50, 75, or 130 can now perform a number of Dynamic Host Configuration Protocol (DHCP) services, including:

- DHCP Server functions, responding to DHCP requests for up to 43 clients at any given time. DHCP server responses provide an IP address and subnet mask. Two address pools of up to 20 IP addresses each can be defined.

  Additionally, up to three hosts, identified by their MAC (Ethernet) addresses, can have an IP address reserved for their exclusive use.
- Managing Plug and Play requests for TCP/IP configuration settings from computers using Microsoft Windows 95 or Windows NT.

- DHCP Spoofing responses, supplying a temporary IP address for a single host. The IP address supplied is always one greater than that of the Pipeline. The IP address is good for only 60 seconds—just long enough to allow a security-card user to acquire the current password from an ACE or SAFEWORD server and bring up an authenticated dial-up session. Once the dial-up session is established, an official IP address can be retrieved from a remote DHCP or BOOTP server.

  This, together with network address translation (NAT), allows a single computer to connect to a remote network that assigns IP addresses dynamically.

How IP addresses are assigned

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

- When the plug-and-play option is enabled (DHCP PNP Enabled=Yes), the Pipeline takes its own IP address, increments it by one, and returns it in the BOOTP reply message along with IP addresses for the Default Gateway and Domain Name Server. Plug-and-play works with Microsoft Windows 95 (and potentially other IP stacks) to assign an IP address and other wide-area networking settings to a requesting device automatically. With plug-and-play you can use the Pipeline to respond to distant networks without having to configure an IP address first.
- If there is an IP address that is reserved for the host, the Pipeline assigns the reserved address.
- If the host is renewing the address it currently has, the Pipeline assigns the host the same address.

  When a host gets a dynamically assigned IP address from one of the address pools, it periodically renews the lease on the address until it has finished using it, as defined by the DHCP protocol. If the host renews the address before its lease expires, the Pipeline always provides the same address.
- If the host is making a new request and there is no IP address reserved for the host, the Pipeline assigns the next available address from its address pools.

  Up to two 20-address pools of contiguous IP addresses are drawn from. Addresses are assigned using the first available address from the first pool or, if there no available addresses in that pool and there is a second pool, the first available address in the second pool.

Configuring DHCP services

To configure a DHCP service, open the following menu:

Ethernet > Mod Config > DHCP Spoofing

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

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

```
20-A00 Mod Config
 DHCP Spoofing...
  DHCP Spoofing=Yes
  DHCP PNP Enabled=Yes
  Renewal Time=10
  Become Def. Router=No
  Dial If link down=No
  Always Spoof=Yes
  Validate IP=Yes
  Maximum no reply wait=5
  IP group 1=181.100.100.100/16
  Group 1 count=1
  IP group 2=0.0.0.0/0
  Group 2 count=0
  Host 1 IP=181.100.100.120
  Host 1 Enet=0080c75Be95e
  Host 2 IP=0.0.0.0/0
  Host 2 Enet=000000000000
  Host 3 IP=0.0.0.0/0
  Host 3 Enet=000000000000
```

1. Set the DHCP Spoofing parameter to Yes to enable any DHCP service. This parameter, which was included in earlier versions of the Ascend software, now has a different meaning. It must be Yes for any DHCP service to be enabled. If it is set to No, other settings in this menu are ignored.
2. Set the DHCP PNP Enabled parameter to Yes to enable Plug and Plug. Setting this parameter to Yes and DHCP Spoofing set to Yes is all that is required to enable Plug and Play support.

3. Renewal Time specifies how long a DHCP IP address lives before it needs to be renewed. It applies to DHCP spoofed addresses and DHCP server replies. If the host renews the address before it expires, the Pipeline provides the same address. Plug and Play addresses always expire in 60 seconds.
4. Become Default Router is an option you can set to advertise the address of your Pipeline as the default router for all DHCP request packets.
5. Dial If Link is Down is used with DHCP spoofing in conjunction with BOOTP Relay. Previously DHCP spoofing had to be disabled for BOOTP relay to work because the two functions attempted to respond to the same requests in different ways. Now if both features are enabled, and no wide area network links are active, the Pipeline performs DHCP spoofing. As soon as the dialed link is established, the Pipeline stops spoofing and acts as a BOOTP relay agent.
6. Set Always Spoof as follows:
   - **Yes enables the DHCP server.** A DHCP server always supplies an IP address for every request, until all IP addresses are exhausted.
   - **No enables DHCP spoofing.** DHCP spoofing only supplies an IP address for a single host on the network. It does not respond to all requests.

   If both DHCP Spoofing and Always Spoof are Yes, the DHCP *server* feature is enabled. If DHCP Spoofing is Yes and Always Spoof is No, DHCP *spoofing* is enabled and works as it did in earlier releases when the value of Always Spoof was Yes.
7. Set Validate IP to Yes to check if a spoofed address that is about to be assigned is already in use, and if it is, automatically assign another address.
8. Set Maximum No-Reply Wait only if you are validating IP addresses. To validate the IP address, DHCP sends an ICMP echo (ping) to check if the address is in use. The maximum time it waits for a reply is determined by this setting. The default is 10 seconds.
9. To assign IP addresses dynamically, set the IP Group 1 parameter to the first address for the IP address pool.
10. Set the Group 1 Count parameter to the number of addresses in the pool. The pool can contain up to 20 addresses.
11. To define an additional address pool for dynamic address assignment, set the IP Group 2 parameter to the first address for the second IP address pool.

12 Set the Group 2 Count parameter to the number of addresses in the pool. The second pool, which can also contain up to 20 addresses, is used only if there are no addresses available in the first pool.

13 To reserve an IP address for a particular host, set the Host 1 IP parameter to the IP address for the host.

14 Set the Host 1 Enet parameter to the MAC (Ethernet) address of the host. The MAC address is normally the Ethernet address of the network interface card that the host uses to connect to the local-area network. The DHCP server assigns this host the IP address you specify whenever it gets a DHCP request for an IP address from the host with that MAC address.

15 To reserve an IP address for another host, set the Host 2 IP parameter to the IP address for the host.

16 Set the Host 2 Enet parameter to the MAC (Ethernet) address of the host.

17 To reserve an IP address for another host, set the Host 3 IP parameter to the IP address for the host.

18 Set the Host 3 Enet parameter to the MAC (Ethernet) address of the host.

Setting up a DHCP server

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

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

Additionally, you might set these parameters:

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

Setting up Plug and Play support

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

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

Setting up DHCP spoofing

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

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

Additionally, you might set these parameters:

```
  Renewal Time=nn
  Become Def. Router=Yes|No
  Dial If Link Down=Yes|No
  Validate IP=Yes
  Maximum no reply wait=n
```

Parameter reference

DHCP Spoofing

Description: Enables or disables all of the DHCP features.

Usage: Press Enter to cycle through the choices.

- Yes enables all DHCP features.
- No disables all DHCP features.

  Yes is the default.

Parameter Location: Ethernet > Mod Config > DHCP Spoofing

See Also: Always Spoof

DHCP PNP Enabled

Parameter Description: Determines whether the Pipeline will automatically assign an IP address, and return it along with the Default Gateway and Domain Name Server IP addresses to the requesting device on a remote network. The default is Yes.

Usage: Press Enter to toggle between Yes (the default) and No.

Parameter Location: Ethernet > Mod Config > DHCP Spoofing

See also: BOOTP Relay

Renewal Time

Description: Specifies the lease time for a dynamically assigned IP address. This is the time in which the host is assigned the IP address, as defined by the DHCP protocol. If the host renews the address before its lease period expires, the DHCP service reassigns the same address.

Usage: Enter a length of time in seconds. The default is 10.

Parameter Location: Ethernet > Mod Config > DHCP Spoofing

Become Default Router

Description: Determines whether the Pipeline should advertise itself as the default router in DHCP responses.

Usage: Press Enter to toggle between choices.

- Yes indicates that the Pipeline performing DHCP responses is the default router.
- No does not advertise the Pipeline as the default. No is the default.

Parameter Location: Ethernet > Mod Config > DHCP Spoofing

See Also: BOOTP Relay Enable

Dial If Link Down

Description: When the WAN link is down, determines whether the Pipeline should dial the first Connection Profile to send a DHCP reply.

Usage: Press Enter to toggle between choices.

- Yes forces the first Connection Profile to always be dialed (when the WAN link is down) whenever a DHCP client request is responded to.
- No lets the Pipeline connect according to settings already in place in the environment, such as according to the current TCP/IP settings, or settings for any other network management software in use.
  No is the default.

Parameter Location: Ethernet > Mod Config > DHCP Spoofing

See Also: See also: BOOTP Relay Enable

Always Spoof

Description: Determines how the Pipeline responds to DHCP requests:

- It can be a DHCP server for up to 43 hosts and assign addresses from its own address pools.
- It can perform DHCP spoofing for a single host by providing a temporary IP address just long enough for a DHCP server on the remote network to provide an official address.

When a Pipeline performs DHCP spoofing, it responds to DHCP requests from only one host. It ignores requests from any host other than the first one to send a request.

Usage: Press Enter to cycle through the choices:

- Yes causes the Pipeline to be a DHCP server.
- No enables DHCP spoofing. No is the default.

Dependencies: If DHCP Spoofing is No, this parameter is N/A.

Parameter Location: Ethernet > Mod Config > DHCP Spoofing

See Also: DHCP Spoofing

Validate IP

Description: When a Pipeline receives a DHCP message requesting an IP address, this parameter determines whether the Pipeline checks to see if the address is already in use. If it is, the Pipeline assigns another address.

Usage: Press Enter to cycle through the choices:

- Yes enables validation of IP addresses.
- No disables validation of IP addresses.

  No is the default.

Dependencies: If DHCP Spoofing and Always Spoof are not both Yes, this parameter is N/A.

Parameter Location: Ethernet > Mod Config > DHCP Spoofing

See Also: DHCP Spoofing, Always Spoof

Maximum No Reply Wait

Description: When a Pipeline handles a DHCP message that requests an IP address and the value of the Validate IP parameter is Yes, it sends an ICMP echo (ping) message to check if the address is already in use. This parameter specifies a maximum duration, in seconds, for two actions related to this check:

- It specifies the length of time during which the Pipeline waits for a response to the ICMP echo message. If the Pipeline does not receive a response during this interval, it assumes that the address is not being used and reserves the address for the host requesting it.

  Note: During the time the Pipeline is validating the address, it ignores the original DHCP request and any subsequent requests from the same host. The host continues to send DHCP requests, however, as specified in the DHCP protocol.
- Once the Pipeline has determined that the address is available, it assigns the host the address if it receives another DHCP request from the host within the number of seconds specified by this parameter. If the Pipeline does not receive the DHCP request during this interval, the Pipeline stops reserving the address.

Usage: Press Enter to open a text field and enter a number between 5 and 300.

10 is the default.

Press Enter to close the text field.

Dependencies: If the DHCP Spoofing and Always Spoof parameters are not both Yes, this parameter is N/A. If Validate IP is No, the Pipeline does not validate the addresses it assigns, regardless of the value of this parameter.

Parameter Location: Ethernet > Mod Config > DHCP Spoofing

See Also: DHCP Spoofing, Always Spoof, Validate IP

IP Group 1

Description: The meaning of this parameter depends on whether the Pipeline is configured to be a DHCP server (when both DHCP Spoofing and Always Spoof are Yes) or is configured to perform DHCP spoofing (when DHCP Spoofing is Yes and Always Spoof is No):

- If the Pipeline is configured to be a DHCP server, this is the address and subnet mask for the first IP address in a pool of addresses used for dynamic address assignment.
- If the Pipeline performs DHCP spoofing, this parameter specifies a spoof address: a temporary address that is provided to the host while the actual IP address is obtained from a DHCP server on the remote network.

Usage: Press Enter to open a text field and enter the IP address and subnet mask.

The address consists of four numbers between 0 and 255, separated by periods. Separate the subnet mask from the address with a slash. To specify the first address in the pool, the IP address must be a valid IP address on the local Ethernet network. To disable address assignment from this pool, enter `0.0.0.0/0`.

The default value is `192.0.2.1/24`.

Press Enter to close the text field.

```
Example:  10.2.1.1/24
```

In this example, `10.2.1.1` is the IP address. The number 24 represents the number of bits in the subnet mask. Masking 24 bits provides a subnet of `10.2.1.0`.

Dependencies: If DHCP Spoofing is No, this parameter is N/A. The Group 1 Count parameter specifies the number of addresses in the pool. All the addresses in the pool must be on the same subnet, and the subnet must be on the local network. If this parameter is `0.0.0.0/0`, which disables address assignment from this pool, the Group 1 Count parameter must be 0.

Parameter Location: Ethernet > Mod Config > DHCP Spoofing

See Also: DHCP Spoofing, Always Spoof, Group 1 Count, IP Group 2

Group 1 Count

Description: If the Pipeline is configured to be a DHCP server, this parameter determines the number of contiguous IP addresses in the first address pool.

Usage: Press Enter to open a text field and enter number between 0 and 20.

Enter 0 if the IP Group 1 parameter is `0.0.0.0/0` (which disables address assignment from the pool) or if the IP Group 1 parameter specifies a DHCP spoof address.

The default is 1.

Press Enter to close the text field.

Dependencies: If the DHCP Spoofing and Always Spoof parameters are not both Yes, this parameter is N/A. The IP Group 1 parameter specifies the first address in the pool. All the addresses in the pool must be on the same subnet, and the subnet must be on the local network. If you are specifying a pool, the value cannot be 0.

Parameter Location: Ethernet > Mod Config > DHCP Spoofing

See Also: DHCP Spoofing, Always Spoof, IP Group 1

IP Group 2

Description: If the Pipeline is configured to be a DHCP server, this is the address and subnet mask for the first IP address in the second pool of addresses used for dynamic address assignment. A second pool is optional; you need it only if you need to assign more than 20 IP addresses or if you need up to 20 but not

enough contiguous addresses are available. Addresses in the second pool are used only if there are no addresses available in the first pool.

Usage: Press Enter to open a text field and then type the IP address and subnet mask.

The address consists of four numbers between 0 and 255, separated by periods. Separate the subnet mask from the address with a slash. To specify the first address in the pool, the IP address must be a valid IP address on the local Ethernet network. To disable address assignment from this pool, enter `0.0.0.0/0`.

The default value is `0.0.0.0/0`.

Press Enter to close the text field.

`Example: 10.2.1.21/24`

In this example, `10.2.1.21` is the IP address. The number 24 represents the number of bits in the subnet mask. Masking 24 bits provides a subnet of `10.2.1.0`.

Dependencies: If DHCP Spoofing is No, this parameter is N/A. The Group 2 Count parameter specifies the number of addresses in the pool. All the addresses in the pool must be on the same subnet, and the subnet must be on the local network. If this parameter is `0.0.0.0/0`, which disables address assignment from this pool, the Group 2 Count parameter must be 0.

Parameter Location: Ethernet > Mod Config > DHCP Spoofing

See Also: DHCP Spoofing, Always Spoof, IP Group 1, Group 2 Count

Group 2 Count

Description: If the Pipeline is configured to be a DHCP server, this parameter determines the number of contiguous IP addresses in the second address pool.

Usage: Press Enter to open a text field and then type a number between 0 and 20.

If the value is 0, the pool is unavailable.

The default is 0.

Press Enter to close the text field.

Dependencies: If the DHCP Spoofing and Always Spoof parameters are not both Yes, this parameter is N/A. The IP Group 2 parameter specifies the first address in the pool. All the addresses in the pool must be on the same subnet, and the subnet must be on the local network.

Parameter Location: Ethernet > Mod Config > DHCP Spoofing

See Also: DHCP Spoofing, Always Spoof, IP Group 1

Host 1 IP

Description: If the Pipeline is configured to be a DHCP server, this parameter reserves an IP address for the host whose MAC (Ethernet) address is specified by the Host 1 Enet parameter. When the host sends a DHCP message requesting an IP address, the Pipeline always assigns this address.

Usage: Press Enter to open a text field and then type the IP address and subnet mask for the host.

The address consists of four numbers between 0 and 255, separated by periods. Separate the subnet mask from the address with a slash. To assign an address, the IP address must be a valid IP address on the local Ethernet network. To disable address assignment, enter `0.0.0.0/0`.

The default value is `0.0.0.0/0`.

Press Enter to close the text field.

```
Example:  10.2.1.41/24
```

Dependencies: If the DHCP Spoofing and Always Spoof parameters are not both Yes, this parameter is N/A. If you enter a value other than `0.0.0.0/0` for this parameter, you must enter a valid MAC address for the Host 1 Enet parameter. If you disable address assignment by entering `0.0.0.0/0` for this parameter, you must set the Host 1 Enet parameter to `000000000000`.

Parameter Location: Ethernet > Mod Config > DHCP Spoofing

See Also: DHCP Spoofing, Always Spoof, Host 1 Enet

Host 1 Enet

Description: If the Pipeline is configured to be a DHCP server, this parameter specifies a host on the local network for which an IP address is reserved. The reserved address is specified by the Host 1 IP parameter. When the host sends a DHCP message requesting an IP address, it always receives this address.

Usage: Press Enter to open a text field.

To specify a host to be assigned an IP address, type the MAC address of the host's Ethernet interface. To disable address assignment, enter `000000000000`.

The default value is `000000000000`.

Press Enter to close the text field.

```
Example:  00d07b5e16e3
```

Dependencies: If the DHCP Spoofing and Always Spoof parameters are not both Yes, this parameter is N/A. If you enter a value other than `000000000000` for this parameter, you must enter a valid IP address for the Host 1 IP parameter. If you disable address assignment by entering `000000000000` for this parameter, you must set the Host 1 IP parameter to `0.0.0.0/0`.

Parameter Location: Ethernet > Mod Config > DHCP Spoofing

See Also: DHCP Spoofing, Always Spoof, Host 1 IP

Host 2 IP

Description: If the Pipeline is configured to be a DHCP server, this parameter reserves an IP address for the host whose MAC (Ethernet) address is specified by the Host 2 Enet parameter. When the host sends a DHCP message requesting an IP address, the Pipeline always assigns this address.

Usage: Press Enter to open a text field and then type the IP address and subnet mask.

The address consists of four numbers between 0 and 255, separated by periods. Separate the subnet mask from the address with a slash. The IP address must be a valid IP address on the local Ethernet network.

The default value is `0.0.0.0/0`.

Press Enter to close the text field.

`Example: 10.2.1.42/24`

Dependencies: If the DHCP Spoofing and Always Spoof parameters are not both Yes, this parameter is N/A. If you enter a value other than `0.0.0.0/0` for this parameter, you must enter a valid MAC address for the Host 2 Enet parameter. If you disable address assignment by entering `0.0.0.0/0` for this parameter, you must set the Host 2 Enet parameter to `000000000000`.

Parameter Location: Ethernet > Mod Config > DHCP Spoofing

See Also: DHCP Spoofing, Always Spoof, Host 2 Enet

Host 2 Enet

Description: If the Pipeline is configured to be a DHCP server, this parameter specifies a host on the local network for which an IP address is reserved. The reserved address is specified by the Host 2 IP parameter. When the host sends a DHCP message requesting an IP address, it always receives this address.

Usage: Press Enter to open a text field.

To specify a host to be assigned an IP address, type the MAC address of the host's Ethernet interface. To disable address assignment, enter `000000000000`.

The default value is `000000000000`.

Press Enter to close the text field.

`Example: 00d07b5e16e4`

Dependencies: If the DHCP Spoofing and Always Spoof parameters are not both Yes, this parameter is N/A. If you enter a value other than `000000000000` for this parameter, you must enter a valid IP address for the Host 2 IP parameter. If you disable address assignment by entering `000000000000` for this parameter, you must set the Host 2 IP parameter to `0.0.0.0/0`.

Parameter Location: Ethernet > Mod Config > DHCP Spoofing

See Also: DHCP Spoofing, Always Spoof, Host 2 IP

Host 3 IP

Description: If the Pipeline is configured to be a DHCP server, this parameter reserves an IP address for the host whose MAC (Ethernet) address is specified by the Host 3 Enet parameter. When the host sends a DHCP message requesting an IP address, the Pipeline always assigns this address.

Usage: Press Enter to open a text field and enter the IP address and subnet mask.

The address consists of four numbers between 0 and 255, separated by periods. Separate the subnet mask from the address with a slash. The IP address must be a valid IP address on the local Ethernet network.

The default value is `0.0.0.0/0`.

Press Enter to close the text field.

```
Example:  10.2.1.43/24
```

Dependencies: If the DHCP Spoofing and Always Spoof parameters are not both Yes, this parameter is N/A. If you enter a value other than `0.0.0.0/0` for this parameter, you must enter a valid MAC address for the Host 3 Enet parameter. If you disable address assignment by entering `0.0.0.0/0` for this parameter, you must set the Host 3 Enet parameter to `000000000000`.

Parameter Location: Ethernet > Mod Config > DHCP Spoofing

See Also: DHCP Spoofing, Always Spoof, Host 3 Enet

Host 3 Enet

Description: If the Pipeline is configured to be a DHCP server, this parameter specifies a host on the local network for which an IP address is reserved. The reserved address is specified by the Host 3 IP parameter. When the host sends a DHCP message requesting an IP address, it always receives this address.

Usage: Press Enter to open a text field.

To specify a host to be assigned an IP address, type the MAC address of the host's Ethernet interface. To disable address assignment, enter `000000000000`.

The default value is `000000000000`.

Press Enter to close the text field.

`Example: 00d07b5e16e5`

Dependencies: If the DHCP Spoofing and Always Spoof parameters are not both Yes, this parameter is N/A. If you enter a value other than `000000000000` for this parameter, you must enter a valid IP address for the Host 3 IP parameter. If you disable address assignment by entering `000000000000` for this parameter, you must set the Host 3 IP parameter to `0.0.0.0/0`.

Parameter Location: Ethernet > Mod Config > DHCP Spoofing

See Also: DHCP Spoofing, Always Spoof, Host 3 IP

DNS list size increased

Previously, the number of DNS addresses listed for terminal server logins was limited to six. Now you can configure up to 35 addresses, which is the maximum supported by BSD.

The new List parameters

List Attempt

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

List Size

The new List Size parameter specifies a number of addresses that will be listed. The maximum number is 35. To use this feature, the List Attempt feature must be enabled, as shown below:

```
20-A00 Mod Config
DNS...
>Domain Name=abc.com
```

```
Sec Domain Name=Yes
Allow As Client DNS
List Attempt=Yes
List Size=6
Client PRI DNS=0.0.0.0
Client Sec DNS=0.0.0.0
```

List Attempt

Description: Enables or disables the users ability to use alternate DNS hosts.

Usage: Enter Yes to enable or No to disable the List Attempt function.

Dependencies: None.

Parameter Location: Ethernet > Mod Config > DNS

See Also: List Size.

List Size

Description: Specifies a number of DNS addresses that will be made accessible to terminal server users in response to a DNS query. The maximum is 35 because BSD has a limit of 35.

Usage: Press Enter to open a text field, and then specify a number between 0 and 35. The default value is 6.

Dependencies: This parameter is N/A if the List Attempt feature is disabled.

Parameter Location: Ethernet > Mod Config > DNS

See Also: List Attempt

User-definable TCP connection retry timeout

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

Previously, the timeout period was not user-definable, and the timeout value was always 170 seconds, which is longer than some client software waits before timing out. When client software times out, the connection is dropped and no remaining addresses on the DNS list are tried. Then, each time the Pipeline restarted, it attempted the same connection that was previously unsuccessful.

To use the feature, set TCP Timeout value from 1 and 200 seconds so that, if necessary, connections to additional host addresses can be attempted before the client software times out. When the timeout value is reached and no connection was made, the Pipeline tries the next address on the list.

Choosing a value for TCP Timeout:

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

A longer timeout is appropriate if the environment includes servers with

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

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

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

TCP time-out

Description: Specifies the length of time during which a Pipeline will attempt to connect to an IP host in a list provided by a DNS server.

Since the first host on the list may not be available, the timeout should be short enough to allow the Pipeline to go on to the next address on the list before the client software times out.

When set, the parameter applies to all TCP connections initiated from the Pipeline, including Telnet, Rlogin, TCP-Clear, and the TCP portion of DNS queries.

Usage: To set the timeout value, select TCP Timeout and enter the number of seconds the Pipeline should wait to connect to an IP address on the DNS list.

The range of values for TCP timeout is 0 to 200 seconds. This specifies the number of seconds after which the Pipeline will stop attempting to connect to an IP address and will proceed to the next address on the list.

Note: There is a built-in maximum number of connect messages the Pipeline will send to attempt to connect to a remote host. When the Pipeline has sent the maximum number of messages to an address on the DNS list it will stop attempting to make a connection to that address, even if the maximum time set in TCP timeout has not yet elapsed.

The default for TCP timeout is 0. If TCP timeout=0, the Pipeline will retry the connection to the address at increasingly larger intervals until it sends the maximum number of start-connection messages. This takes approximately 170 seconds, but can take longer if the Pipeline is running a large number of other tasks. If the client software times out before the Pipeline makes a connection or proceeds to the next address on the DNS list, the physical connection is dropped.

Dependencies: The List Attempt parameter in the DNS submenu of the Mod Config menu in the Ethernet Profile must be enabled. This permits the Pipeline to attempt a series of IP addresses. Be aware that the List Attempt parameter does not apply if Telnet and Immediate Telnet are both disabled.

Parameter Location: Ethernet > Mod Config

See Also: List Attempt

Dial-in user DNS server assignments

IP addresses for Domain Name Servers (DNS) can now be set for users who dial into the Pipeline via PPP. Previously, the two DNS addresses that were configured on the Pipeline, using the primary (PRI) DNS parameters, were given to all dial-in clients during IPCP negotiation (which is part of PPP).

This feature allows you to give dial-in clients primary and secondary DNS server addresses by supplying information in the parameters PRI DNS and SEC DNS. If DNS servers are not specified for the dial-in user, the Pipeline supplies the IP addresses for the two DNS servers.

DNS information is supplied based on these rules:

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

Configuring DNS servers in the Ethernet profile

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

1 Open the Ethernet > Mod Config > DNS menu.
For example:

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

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

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

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

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

Configuring DNS servers in the Connection profile

To configure DNS servers in the Connection profile:

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

```
30-100 Connections
IP Options...
LAN Adrs=0.0.0.0/0
WAN Adrs=0.0.0.0
IP Adrs=0.0.0.0/0
Metric=7
Preference=100
Private=No
RIP=Off
Pool=0
Multicast Client=No
Multicast Rate Limit=5
```

```
Client Pri DNS=111.11.11.1
Client Sec DNS=111.11.11.2
Client Assign DNS=Yes
```

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

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

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

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

4 Select Yes or No for Client Assign DNS.

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

Parameter reference

Allow as Client DNS

Description: Specifies whether the local DNS servers should be made accessible to PPP connections if the client DNS servers are unavailable.

Client DNS configurations define DNS server addresses that will be presented to WAN connections during IPCP negotiation. They 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, and a connection-specific configuration that applies to that connection only. The global client addresses are used only if none are specified in the Connection profile.

This parameter acts as a flag to enable the Pipeline to present the local DNS servers to the WAN connection when all client DNS servers are not defined or available.

Usage: Specify Yes or No. No is the default.

- Yes allows clients to use the local DNS servers.
- No prevents clients from using the local DNS servers.

Location: Ethernet > Mod Config > DNS

See Also: Client Assign DNS, Client Pri DNS, Client Sec DNS

Client Assign DNS

Description: Specifies whether client DNS server addresses will be presented while this connection is being negotiated.

Usage: Specify Yes (to use client DNS servers) or No. No is the default.

Example: Client Assign DNS = no

Location: Ethernet > Connections > *any profile* > IP Options

See Also: Client Pri DNS, Client Sec DNS

Client Pri DNS

Description: Specifies a primary DNS server address to be sent to any client connecting to the Pipeline. Client DNS has two levels: a global configuration that applies to all PPP connections, and a connection-specific configuration that applies to that connection only. The global client addresses are used only if none are specified in the Connection profile. You can also choose to present your local DNS servers if no client servers are defined or available.

Usage: Specify the IP address of a DNS server to be used for all connections that do not have a DNS server defined. The default value is 0.0.0.0.

Example: Client Pri DNS=10.9.8.7/24

Location: Ethernet > Mod Config > DNS; Ethernet > Connections > *any profile* > IP Options

---

Client Sec DNS

Description: Specifies a secondary DNS server address to be sent to any client connecting to the Pipeline. Client DNS has two levels: a global configuration that applies to all PPP connections, and a connection-specific configuration that

applies to that connection only. The global client addresses are used only if none are specified in the Connection profile. You can also choose to present your local DNS servers if no client servers are defined or available.

Usage: Specify the IP address of a secondary DNS server to be used for all connections that do not have a DNS server defined. The default value is 0.0.0.0.

Example: Client Sec DNS=10.9.8.7/24

Location: Ethernet > Mod Config > DNS; Ethernet > Connections > *any profile* > IP Options

Sec Domain Name

Description: Specifies a secondary domain name that the Pipeline can search using DNS. The Pipeline performs DNS lookups in the domain configured in Domain Name first, and then in the domain configured in Sec Domain Name.

Usage: Specify a secondary domain name. You can enter up to 63 characters.

Example: Sec Domain Name=xyz.com

Location: Ethernet > Mod Config > DNS

Local DNS host address table option added

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

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

You can also specify that the local DNS table is automatically updated when a connection to a host whose name matches one in the local DNS table is successfully resolved by the remote DNS. When the table is updated, the returned IP

address list from the remote server replaces the stored IP addresses for that host name in the local DNS list.

You can check the list of host names and IP addresses in the table using the term-serv command Show Dnstab.

Configuring the local DNS table

To enable and configure the local DNS table:

1 Display Ethernet Profile: Ethernet > Mod Config > DNS menu.

2 Select List Attempt=Yes to allow a list of the IP addresses to be displayed when using the terminal server command Dnstab Entry.

3 Select List Size and enter the number of entries you want in the list.

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

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

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

For example:

- If you set List Size=4 and the remote DNS returns 3 entries, the entire list of IP addresses in the local DNS table is cleared and the three returned addresses are entered for the entry.
- If the local DNS table already contains 35 IP addresses for an entry and the remote DNS server returns only 4, or if you set List Size=4, the first four IP addresses are entered into the table for the entry and the remaining addresses in the list are set to zero.
- If you set List Size=1, the list can contain only one IP address; any others returned by the remote DNS are ignored. If you change the List Size parameter value from a number greater than one to one, only the first IP address is retained; all others are set to zero the next time the table entry for that name is updated.

4 Select Enable Local DNS Table=Yes.

The default is No.

5 Select Loc DNS Tab Auto Update=Yes to enable automatic updating.

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

Creating the local DNS table

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

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

1 From the terminal server, enter:

```
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 Return.

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 `0` and press Return when the editor prompts you for another entry.

Editing the local DNS table

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

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

1 From the terminal server, enter:

```
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 Return

- 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 `0` 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:
   `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.

Restrictions for names in the local DNS table

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

Show command changes

Additional terminal server Show commands have been added to help you view and edit the DNS table:

- `show ?` displays a list that includes `Dnstab` help.
- `show dnstab` displays the local DNS table.
- `show dnstab ?` displays help for the dnstab editor.

Terminal server dnstab command

The terminal server `dnstab` command has three variations:

dnstab	Displays help information about the DNS table.
dnstab show	Displays the local DNS table.

dnstab entry *n*	Displays a list for entry *n* in the local DNS table. The list displayed includes the entry and all the IP addresses stored for that entry up to a maximum number of entries specified in the List Size parameter. If List Attempt=No, no list is displayed.

Enable Local DNS Table

Description: Enables the use of a local DNS table to provide a list of IP addresses for a specific host when the remote DNS server fails to resolve the host name. The Local DNS table will provide the list of IP addresses only if the host name for the attempted connection matches a host name in the local DNS table.

Usage: Select Yes to enable the local DNS table. No disables the feature.

Parameter Location: Ethernet > Mod Config > DNS

Loc. DNS Tab Auto Update

Description: Enables or disables automatic updating. When automatic updating is enabled, the list of IP addresses for each entry is replaced with a list from the remote DNS when the remote DNS successfully resolves a connection to a host named on the table.

Usage: Select Yes or No to enable or disable updating of the IP addresses.

Dependencies: Enable Local DNS Table must be set to Yes. To display the list of IP addresses for a DNS table entry, List Attempt must be set to Yes, and a value of 1-35 specified for List Size. If you set List Attempt=No, the `dnstab` show command displays only the first IP address on the list.

Parameter Location: Ethernet > Mod Config > DNS

Spoof Adr parameter allows any subnet address and mask

Previously, you were required to specify an ip address in the Spoof Adr parameter with a subnet address and subnet mask that matched the address and mask configured for the MAX. With this release, you can specify a subnet address and subnet netmask different from those on the MAX.

SPID 1

Description: This parameter specifies the ISDN BRI Service Profile Identifier (SPID) associated with My Num A. An SPID is a number assigned to a domestic ISDN BRI line for service identification at the ISDN service provider's central office. It is typically formed by adding a code to the phone number assigned to the line. Your carrier provides you with one or more SPIDs.

All U.S. domestic switch types, except AT&T Point-To-Point, can have two phone numbers. The primary phone number (My Num A) requires a matching primary SPID (SPID 1). The secondary phone number (My Num B) requires a matching secondary SPID (SPID 2).

When you use AT&T Point-to-Point service, only one phone number is assigned to the ISDN BRI line, and no SPIDs are used.

Usage: Press Enter to open a text field. Then, type up to 16 characters; you must limit those characters to numbers, hyphens, and parentheses. The default value is 0 (zero). Press Enter again to close the text field.

Dependencies: Keep this additional information in mind:

- You must enter a value for SPID 1 unless you are using AT&T Point-To-Point (Link Type=P-T-P) or you are operating outside of the U.S.
- If the Pipeline uses only one channel of a multipoint ISDN BRI line and another device uses the other channel, you can choose to operate in single-terminal mode.

  Set one channel to unused (Chan Usage=Unused/Switch or Chan Usage=Switch/Unused), and enter only one SPID. The device sharing the line must enter the other assigned SPID.

- The Pipeline appends the value of SPID 1 with a TID if you are connected to a Northern Telecom switch running NI-1 (Switch Type=NI-1).

Parameter Location: Configure Profile

See Also: Chan Usage, My Num A, My Num B, Sec Num, SPID 2, Switch Type

---

SPID 2

Description: This parameter specifies the ISDN BRI Service Profile Identifier (SPID) associated with My Num B. An SPID is a number assigned to a domestic ISDN BRI line for service identification at the central office (CO). It is typically formed by adding a code to the phone number assigned to the line. Your carrier provides you with one or more SPIDs.

All U.S. domestic switch types, except AT&T Point-To-Point, can have two phone numbers. The primary phone number (My Num A) requires a matching primary SPID (SPID 1). The secondary phone number (My Num B) requires a matching secondary SPID (SPID 2).

When you use AT&T Point-to-Point service, only one phone number is assigned to the ISDN BRI line, and no SPIDs are in use.

Usage: Press Enter to open a text field. Then, type up to 16 characters; you must limit those characters to numbers, hyphens, and parentheses. The default value is 0 (zero). Press Enter again to close the text field.

Dependencies: Keep this additional information in mind:

- You must enter a value for SPID 1 unless you are using AT&T Point-To-Point (Link Type=P-T-P) or you are operating outside of the U.S.
- If the Pipeline uses only one channel of a multipoint ISDN BRI line and another device uses the other channel, you can choose to operate in single-terminal mode.

  Set one channel to unused (Chan Usage=Unused/Switch or Chan Usage=Switch/Unused), and enter only one SPID. The device sharing the line must enter the other assigned SPID.
- The Pipeline appends the value of SPID 1 with a TID if you are connected to a Northern Telecom switch running NI-1 (Switch Type=NI-1).

Parameter Location: Configure Profile

Chan Usage, My Num A, My Num B, Sec Num, SPID 1, Switch Type

Spoof Adr

Description: This parameter specifies an IP address and netmask that will be assigned to the DHCP client when spoofing occurs.

Usage: Press Enter to open a text field, and then type a valid IP address and subnet mask. Press Enter again to close the text field.

Example: 10.0.0.1/24

Parameter Location: Ethernet Profile > Mod Config > DHCP Spoofing.

Dependencies: The Spoof Adr parameter applies only if the DHCP Spoofing and Renewal Time parameters are configured.

See Also: DHCP Spoofing, Renewal Time

IPX Routing

5

Overview

The following new features might affect the way you set up IPX routing on your unit:

IPX Type 20 packet propagation support

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

Since the Pipeline cannot support these types of applications, a parameter has been added to turn off IPX Type 20 packet propagation. You can change the setting to Yes if required.

Handle IPX Type20

Description: Enables or prevents IPX Type 20 packet propagation.

Usage: Press Enter to select Yes, which allows IPX Type 20 packet propagation; or No, which prevents IPX Type 20 packet propagation.

The Pipeline now supports a flag to switch IPX Type 20 propagation ON and OFF.

Parameter Location: Configure > Ethernet > Mod Config > Ether options > IPX SAP Filter

Dependencies: You must have IPX routing enabled, and IPX SAP Filter

New limit for server and route entries

On a large IPX network, the Pipeline does not function properly with more than 300 server and route entries. In order to keep the Pipeline operational with IPX enabled on a large network, this release enforces a maximum limit of 300 server and route entries.

This enhancement includes limit checking for both server and route entries. When the Pipeline reaches its limit of 300, it drops all IPX route and SAP packets containing additional routes and services. This limit results in an incomplete network map, so you need to activate a size-limiting feature, such as proxy SAP or IPX filtering.

Two commands let you see how the Pipeline checks limits for server and route

entries:

```
ipxservinfo
```

```
ipxroutinfo
```

When you enter the Ipxservinfo command in the diagnostic monitor, the following information is displayed:

```
ipx server table info
```

```
ipxservcnt is 20/* IPX server table count */
```

```
ipxservmax is 300/* IPX maximum server table limit */
```

When you enter the Ipxroutinfo command in the diagnostic monitor, the following information is displayed:

```
ipx route table info
```

```
ipxroutcnt is 20/* IPX route table count */
```

```
ipxroutmax is 300/* IPX maximum route table limit */
```

Increase default IPX SAP proxy servers

Some networks are designed to prevent the propagation of RIP and SAP packets from a MAX to a Pipeline. In previous releases, the IPX SAP proxy feature let you point to only one IPX SAP proxy server. If that proxy server was unavailable, remote users could not connect to the network. Now, there are three default IPX SAP proxy servers in the Ethernet profile

IPX SAP Proxy Net#*n*

Description: Specifies a default IPX SAP proxy server (from 1 to 3).

Usage: For each parameter, specify the IPX network number of the server providing the SAP proxy. The default value is 0 (zero).

The Pipeline first attempts to use the server specified by IPX SAP Proxy Net#1. If that server is unavailable, the Pipeline then attempts to use the server specified by IPX SAP Proxy Net#2. If that server is also unavailable, the Pipeline attempts to use the server specified by IPX SAP Proxy Net#3.

Dependencies: If IPX SAP Proxy=No, the IPX SAP Proxy Net#*n* parameter does not apply.

Parameter Location: Ethernet Profile: Ethernet > Mod Config > Ether Options.

Support for IPX without defining an IPX server

You can now specify a route to a destination IPX network without defining an IPX server in the IPX Routes submenu of the Ethernet configuration profile. Previously, if you specified a route without also specifying an IPX server, the Pipeline would put a NULL entry in the SAP table. This feature modifies this behavior so that no entry is placed in the SAP table.

Reaching an IPX network via the network number

There are no user interface changes resulting from this feature. The IPX Routes submenu of the remains the same. You can reach an IPX network by entering the Network number (for example, Network=00123456) without specifying the Server Name and Server Type.

Optimized access for dial-in NetWare clients

In previous releases, the Pipeline assumed that the far end of an incoming IPX connection was another IPX router. After answering the call, the Pipeline could recognize the caller as a client via the Peer=Dialin setting in the caller's Connection profile. For dial-in Windows 95 clients with no configured profile, the connection could take more than a minute to establish and then the client could not see NetWare servers on the local network.

Now the Answer profile also contains a Peer parameter to enable the Pipeline to treat incoming IPX connections as clients even when configured profiles are not in use.

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

The Dialin setting tells the Pipeline to negotiate inbound IPX calls as if the far end is a dial-in NetWare client.

The following list shows the new Peer parameter as well as other required parameters with example values:

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

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

Required settings

To use this feature, configure the Pipeline as follows:

- Calls for which no Connection profile is found must be answered.

  The call might require authentication, or use SecureID passwords.

  The dial-in client must be running PPP software.

- IPX routing must be enabled in the PPP Options submenu of the Answer profile, and the IPX network number of the router's Ethernet interface must be configured in the Ethernet profile.

- Specify an IPX Pool number in the Ethernet profile, so that the Pipeline can route to dial-in clients.

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

  Note: The Pipeline does not send RIP and SAP advertisements across the connection and ignores RIP and SAP advertisements received from the far end. However, it does respond to RIP and SAP queries received from dial-in clients.

IPX filters

IPX packets can now be filtered using an IPX filter interface. In previous releases, filters supported only IP and Generic interfaces. A new selection has been added that provides parameters for filtering IPX packets.

A new IPX filter type can now be specified, for example:

```
Filter name
    In filter 01
    Valid=Yes
    Type=IPX
    Generic...
    Ip...
    Ipx...
```

When the IPX filter type is specified, the following IPX submenu is available:

```
Ipx...
    Forward=No
    Src Network Adrs=cfff0000
    Dst Network Adrs=cf088888
    Src Node Adrs=111222333
    Dst Node Adrs=aaabbbccc
    Src Socket Cmp=equal
    Src Socket #=0451
    Dst Socket Cmp=equal
    Dst Socket #=0015
```

The Forward parameter works just as it does for other filter types. If it is set to No, a matching packet is discarded. The following new filter parameters are supported:

- Src Network Adrs

  The source IPX network address. Either the source or destination address (or both) must be specified.

- Dst Network Adrs

  The destination IPX network address. Either the source or destination address (or both) must be specified.

- Src Node Adrs

A valid IPX node address. The node address ffffffffffff means all nodes in the specified source network. This value must be specified if the Src Network Adrs is not null.

- Dst Node Adrs

  A valid IPX node address. The node address ffffffffffff means all nodes in the specified destination network. This value must be specified if the Dst Network Adrs is not null.

- Src Socket Cmp and Src Socket #

  Some NetWare services communicate across specific sockets; for example, file servers typically use socket 0451. If you specify the source socket number, you can also specify the type of comparison to be made between the source socket for an IPX packet and the value specified in this filter. You can specify that the filter matches the packet if the source socket number is equal, not-equal, less-than, or greater-than the one specified in the filter.

- Dst Socket Cmp and Dst Socket #

  If you specify the destination socket number, you can also specify the type of comparison to be made between the destination socket for an IPX packet and the value specified in this filter. You can specify that the filter matches the packet if the destination socket number is equal, not-equal, less-than, or greater-than the one specified in the filter.

SPX spoofing added for IPX

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

In previous software versions these repeated watchdog packets from the client's SPX watchdog kept the WAN connection up unnecessarily. Now the Pipeline lets Netware SPX clients stay logged in without keeping the WAN connection up in times of inactivity. The Pipeline automatically responds to SPX watchdog requests from the LAN with a look-alike (spoofed) SPX-watchdog-reply packet, and drops any SPX-watchdog keep-alive packets from the LAN, without sending

them on to the WAN.

You do not need to set any parameters.

Note: Routers on both ends of the connection must support this feature for it to function.

Security

6

Overview

The following new features might affect the way you set up security on your unit:

Secure Access support

The Pipeline now supports Secure Access Management (SAM), a graphical user interface for creating IP firewalls.

Refer to the *Ascend Secure Access User's Guide* (part number 7820-0429-001) for complete instructions on using SAM and adding firewalls to your Pipeline.

Ascend currently supports simple "static" packet filters. Each connection may have a call and/or a data filter profile. Each filter profile may have up to 12 inbound and 12 outbound packet filters. Each packet filter may be either a generic or an IP filter.

Secure Access adds these features to the current Ascend filtering:

- Easily write "dynamic" packet filters, also known as "firewalls," which permit traffic to be normally blocked except when triggered by an event, such as an inbound or outbound connection request.
- Log traffic passing through the router and thus provide an audit trail to track IP connections and packet content.
- Send an appropriate ICMP message when a packet is not forwarded due to a firewall.

A limitation of current Ascend packet filters and most other router packet filters is that they are not able to securely deal with a number of IP protocols. Secure Access uses "dynamic" packet filters, which allow the firewall to block all packets except for those that are specifically required for a single session and only for the length of the session.

Unlike the current limitations on the number of filters in a profile, Secure Access does not place a limit on the number of packet filters (that is "rules") in a firewall profile. The only limitation is based on the compressed size of the firewall profile or the limit of memory on the router. Current Ascend packet filters continue to be supported and so no configuration changes on Pipelines are necessary until Secure Access firewalls are actually used.

To enable Secure Access on your Pipeline, you must obtain a hash code.

Using SAM

Configuration of a firewall is done on a network-connected workstation external to the Pipeline using the SAM graphical user interface. You then upload the com-

pleted firewall to the Pipeline where it will be used.

The external application, Secure Access Manager (SAM), allows you to select what services will be permitted to pass through the firewall and what hosts are allowed access to the services.

Once the data has been filled in, SAM can upload the firewall profile to an Pipeline or save the profile into a file.

Refer to the *Ascend Secure Access User's Guide* (part number 7820-0429-001) for complete instructions on using SAM and adding firewalls to your Pipeline.

New menu added to the Telnet interface

A new Firewalls menu has been added to the Telnet interface. For example:

```
20-600 Firewalls
  >20-601 Sales
```

When Secure Access has been enabled on the Pipeline, this menu appears and stores all the firewalls that have been downloaded to your system using SAM. When you open the Firewalls menu, the submenu is listed. For example:

```
20-601 Sales
>Name=Engineering
 Version=1
 Length=2936
```

Note that only the Name field can be edited. The Version and Length parameters are determined by the firewall you create in SAM. Firewalls must be modified using SAM.

Firewall numbers in the Telnet interface

To ensure backward compatibility with the current Ascend filter implementation, you must number firewalls created with SAM differently than filters created using the Telnet interface. You can continue to assign existing filters to Profiles exactly as before. However, if you want to assign a firewall created with SAM to a Profile, you must add 100 to the last two digits of its index in the telnet interface. The numbering scheme for filters is:

- 0 indicates that no filtering is being used

- 1-99 indicates that a filter created using the Telnet interface is being used
- 100-199 indicates that a filter created using SAM is being used.

For example, suppose you have already created these Ascend filters:

```
90-500 Filters
>90-501 IP Call
 90-502 NetWare Call
 90-503 AppleTalk Call
 90-504 Engineering
 90-505 Test Eng
 90-506 Marketing
 90-507
 90-508
 90-500
 90-510
 90-511
 90-512
```

If you want to use the Engineering filter in a Connection or Mod Config Profile, enter the number 4 in the Data Filter or Call Filter field (in a Connection Profile) or in the Filter field (in the Mod Config, Ether options field).

Now suppose you have created a firewall using SAM and downloaded them to your Pipeline. The Firewalls menu may look similar to this:

```
20-600 Firewalls
>20-601 Sales
 20-602
 20-603
 20-604
```

To use the Sales filter in a Connection or Mod Config Profile, enter the number 101 in the Data Filter or Call Filter field (in a Connection Profile) or in the Filter field (in the Mod Config, Ether options field).

Assigning firewalls to a Connection Profile

Firewalls assigned to a Connection Profile are used to filter incoming or outgoing traffic on a WAN connection. Filters assigned to a Connection Profile are activated whenever the WAN session comes online.

To assign a firewall to a Connection profile:

1. Create a firewall filter using SAM.
2. Download it to the Pipeline.
3. Select Ethernet, Connections, *a Connection Profile*, Session options.
4. Enter the number of the firewall filter you want to use in the Data filter field.

   This number is derived from the number in the Firewall menu by adding 100 to the last 2 digits of the firewall index. For example, if the firewall is number 20-503, enter number 103 in the Data Filter field.
5. Exit the Connection Profile and save your changes.

Assigning firewalls to the Mod Config profile

Firewalls assigned to the Mod Config Profile are used to filter incoming or outgoing traffic on the Ethernet interface. Filters assigned to a the Mod Config Profile are activated as soon as you save the changes to the Mod Config Profile.

To assign a firewall to the Mod Config Profile, do the following:

1. Create a firewall filter using SAM.
2. Download it to the Pipeline.
3. Select Ethernet, Mod Config, Ether options.
4. Enter the number of the firewall filter you want to use in the Filter field.

   This number is derived from the number in the Firewall menu. For example, if the firewall is number 20-503, enter number 103 in the Data Filter field.
5. Exit the Connection Profile and save your changes.

New Sys Options field

Secure Access has been added to the Sys Option window. All software that includes the Secure Access feature will include a “Sec Acc” field. If the feature has not yet been enabled, the option will be marked as “Not Inst”. If the feature has been enabled, the option will be marked as “Installed.”

```
00-100 Sys Options
>Switched Installed^
 Frm Rel Installed
 Sec Acc Installed  V
```

New parameters

These new parameters have been added or enhanced to support firewalls on Ascend products:

- Name
- Version
- Length

Name

Description: Specifies the name of the firewall. This name is originally created using the Secure Access Manager (SAM) graphical user interface.

Usage: Press Enter to open a text field. Then, type the name of the firewall. Press Enter again to close the text field.

Parameter Location: Ethernet > Firewalls > *any Firewall*

Version

Description: Each firewall contains a version number to ensure that any firewall that is uploaded to the router will be compatible with the firewall software on the router. Secure Access Manager (SAM) checks the version number before uploading a firewall. In the event that a router with a stored firewall profile receives a code update that make the existing firewall incompatible, a default firewall is enabled, permitting only Telnet access to the Pipeline.

Usage: This parameter cannot be edited.

Parameter Location: Ethernet > Firewalls > *any Firewall*

Length

Description: Specifies the length of the firewall uploaded to the Pipeline from Secure Access Manager (SAM).

Usage: This parameter cannot be edited.

Parameter Location: Ethernet > Firewalls > *any Firewall*

Syslog messages

Syslog messages may be generated for packets seen by the 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 created by firewalls will use the standard 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> <local> <direction> <remote> <length> <frag>
  <log> <tag>
  ```
  – <protocol> is the 4 hexadecimal digit Ether Type, or the network protocol name—"arp," "rarp," "ipx," "appletalk."

    <protocol>, for IP protocols, is either the IP protocol number (up to 3 decimal digits) or one of the following names:

    ip-in-ip

    tcp

    icmp

    udp

    esp

    ah

    In the special case of icmp, it will also include the ICMP Code and Type ([Code]/[Type]/icmp).

  – <local>, 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.

<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]).

– <direction> is an arrow “<-”, “->” showing the direction (receive and send respectively) in which the packet was traveling.

– <remote>, for non-IP protocols, has the same format as <local> non-IP packets but shows the destination Ethernet MAC destination address of transmitted packets and the source Ethernet MAC address of received packets.

  <remote>, 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.

– <length> is the length of the packet in octets (8-bit bytes).

– <frag> is used to report “frag” if the packet has a non-zero IP offset or the IP More-Fragments bit is set in the IP header.

– <log> is used to report one or more messages based upon the packet status or packet header flags. The packet status messages include:

  corrupt—the packet is internally inconsistent

  unreach—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 will only be displayed for the initial packet which has the SYN flag and not the ACK flag set.

– <tag> contains any user defined tags specified in the filter template used by SAM.

Filter persistence

Filter persistence has been added to the Connection Profile of all Pipelines that support Filter Profiles. The Filter Persistence parameter must be set to Yes to allow a connection's firewalls to persist when the connection is torn down, such as by connection timeout. The default case is No, implying that, by default, connection firewalls do not persist when a call is terminated.

Note: Typically a firewall will persist for about an hour after its associated connection has been torn down.

Background on firewall and filter persistence

The idea of filter persistence is intended to allow an Pipeline to preserve its filter/firewall specifications throughout the lifetime of its connections.

Firewalls differ from filters in that firewalls have been designed to alter their behavior as traffic passes through them, where filters remain unchanged through their lifetimes. This has required a change in the way firewalls and filters are associated with connections.

Ascend filters as they were originally implemented provided for the construction and destruction of filters whenever the state of a connection changed. This causes the Pipeline to create and destroy filters during connection state changes without any reference to the state of the filters.

With Secure Access Firewalls, it is necessary to preserve the firewall state across the many transitions that connections may experience. Where filters could be built or destroyed at any time to accommodate changes due to Multilink and idle-inactivity conditions, firewalls cannot.

To resolve this problem, Ascend filters and firewalls can now be persistent. A persistent filter or firewall is maintained even when its associated connection becomes inactive. Additionally, the filter or firewall can be applied when an additional session becomes associated with a connection, as is the case with additional channels of an MPP connection.

Note: Firewalls must have persistence to work correctly, but filters do not.

Filter persistence and Connection Profiles

Connection Profiles describe different contact sites. Perhaps, for a small office, one profile would apply to a corporate home office, and another profile would apply to an Internet service provider. In each case, the Pipeline user would like to use the Secure Access Firewall capability to prevent unauthorized incursions into the local network by others.

With dial-on-demand and automatic call timeout, the dynamic firewall capabilities of Secure Access Firewall would prevent in-progress TCP sessions (such as telnet or rlogin) from proceeding after a call termination and restart (due to inactivity, for example). Without persistence, a new firewall is constructed when a call starts up with no knowledge of any TCP sessions in progress, and consequently would block packets for those sessions when starting the line back up. This has the effect of rendering the in-progress telnet (or rlogin, etc.) sessions inoperative, possibly destroying work in progress that is dependent on them.

Filter persistence is a way to tell the Pipeline to keep a firewall around even after the call is terminated. When a new call is placed to (or is received from) the same station, the Pipeline remembers the original firewall and uses it as if the call had never been terminated. Thus, the user can continue working without loss.

Conversely, there may be times when a single Connection Profile is used for several different sites. This might be the case if you use the same Connection Profile to describe multiple different callers. In this case, you do not want the filters and firewalls to be persistent, since the Pipeline cannot know if calls are arriving from the same users.

Filter persistence

Description: Specifies whether the filter or firewall assigned to a Connection Profile should persist after the call has been disconnected.

Usage: Press Enter to cycle through the choices:

- Yes specifies that the filter or firewall assigned to this Connection Profile will persist after the connection has been torn down.

  Note: Typically a firewall will persist for about an hour after its associated connection has been torn down.

- No specifies that the filter or firewall assigned to this Connection Profile will not persist after the connection has been torn down.

No is the default.

Parameter Location: Ethernet > Connections > *profile* > Session options

See Also: Call Filter, Data Filter, Name, Version, Length

MS-CHAP support

Support for the Microsoft Challenge Handshake Authentication Protocol (MS-CHAP) has been added for Windows NT systems.

You can now configure a Pipeline to send or receive MS-CHAP authentication. MS_CHAP authentication is described in detail at Microsoft's Web site at the following address:

`ftp://ftp.microsoft.com/DEVELOPR/RFC/chapexts.txt`

Note: MS-CHAP with DES and MD4 encryption, supported in this release, works in Windows NT environments only. The Pipeline can authenticate a Windows NT system and a Windows NT system can authenticate a Pipeline.

Configuring a Pipeline for MS-CHAP authentication

The options available for the Recv Auth= parameters in the PPP Options submenu of the Answer profile have been changed. Previously, Recv Auth=Either was available to enable the Pipeline to authenticate received traffic using PAP or CHAP. This option has been changed to Recv Auth=PAP/CHAP/MS-CHAP, which enables the Pipeline to authenticate using any of the authentication protocols (PAP, CHAP, or MS-CHAP) when communicating with Windows NT systems.

To configure a Pipeline to authenticate using MS-CHAP, select one of the following options for Recv Auth= shown in below.

Value for Recv Auth=	Description
EITHER	Allows authentication if the remote peer can authenticate using any of the designated authentication schemes.

Value for Recv Auth=	Description
MS-CHAP	Allows authentication only if the remote peer uses MS-CHAP for authentication.

To configure a Pipeline to send MS-CHAP authentication, select Send Auth=MS-CHAP in the Encaps options submenu of the Connection profile. With this option selected, the Pipeline will only continue authentication if the remote peer also supports MS-CHAP authentication.

Called number authentication supported

This feature adds authentication by Called Number. It is similar to authentication by Caller ID (CLID), but uses the number (ID) of the unit being called instead of the number of the calling unit.

Configuring calling or called number authentication

To support configuration for Called Number Authentication, the following changes have been made:

Profile	Description of change
Connection profile	A new field, Called #, has been added. Called # is usually the same as Dial #, but without the trunk group or dialing prefix prepended.
Answer profile	Clid Auth= is changed to Id Auth=.

Setting	Description of change
Ignore	Ignore calling number or called number.
Prefer	Use Calling number ID authentication, but if it's not available, use name/password authentication.

Setting	Description of change
Require	Use Calling number ID authentication.
Called Require	Same as Require, except use the called number rather than the calling number ID.
Called Prefer	Same as Prefer, except use the called number rather than the calling number ID.

Called #

Description: Adds authentication by Called Number using the number (ID) of the unit being called instead of the number of the calling unit. Called # is the same as Dial #, but without the trunk group or dialing prefix prepended.

Usage: Select the parameter and cycle through the possible choices.

- Ignore. Ignore calling number or called number.
- Prefer. Use Calling number ID authentication, but if it's not available, use name/password authentication.
- Require. Use Calling number ID authentication.
- Fallback does not apply to the Pipeline.
- Called Require. Same as Require, except uses the called number rather than the calling number ID.
- Called Prefer. Same as Prefer, except use the called number rather than the calling number ID.

Dependencies: You should also supply all the information to use name/password authentication in case the called number is blocked (using Caller ID blocking from the phone company). Both types of authentication *are not* performed for the same connection.

Parameter Location: Ethernet > Connections > *any profile*

See Also: Id Auth

Set Disconnect cause code for CLID auth

When Caller ID authentication fails in an ISDN connection, the Pipeline sends a Disconnect message. The Cause Element in the Disconnect message can give an idea of why the CLID authentication failed. You can set the Disconnect cause code for CLID authentication failures to ‘User Busy’ or ‘Normal call clearing’.

Select the Disconnect Cause value in the Ethernet > Mod Config > Auth profile:

```
X0-X00 Mod Config
 Auth...
  CLID Fail Busy=No
  APP Server=No
  APP Host=N/A
  APP Port=N/A
```

CLID Fail Busy

Description: Indicates the Disconnect cause when Called ID authentication fails due to a timeout.

- No sets the Disconnect cause code to ‘Normal call clearing’ and is the default.
- Yes sets the Disconnect cause code to ‘User Busy’.

Usage: Select the parameter and press Enter to cycle through the available settings. Press Esc to exit the parameter.

Dependencies: CLID authentication must be enabled in order to set this parameter. Set it in Ethernet > Answer > ID Auth.

Parameter Location: Ethernet > Mod Config > Auth.

See Also: ID Auth.

Expect callback added to dialout profile

A parameter has been added to the Telco options submenu of the Connections menu for configuring the Pipeline to expect a callback from the machine called.

This prevents problems that arise when CLID is set to Required on the machine that is expected to callback.

How Expect Callback works

When Pipeline initiates a call and the call gets through, the called machine hangs up on the incoming caller and then immediately initiates a call to that destination (callback) before performing password authentication.

For example, in the figure below, ping or telnet is initiated through a MAX to a Pipeline and CLID is set to Required on the Pipeline (the side that will be doing the callback), the Pipeline will reject the incoming call before answering it. To the MAX (the initiating side), it appears as if the call never got through at all.

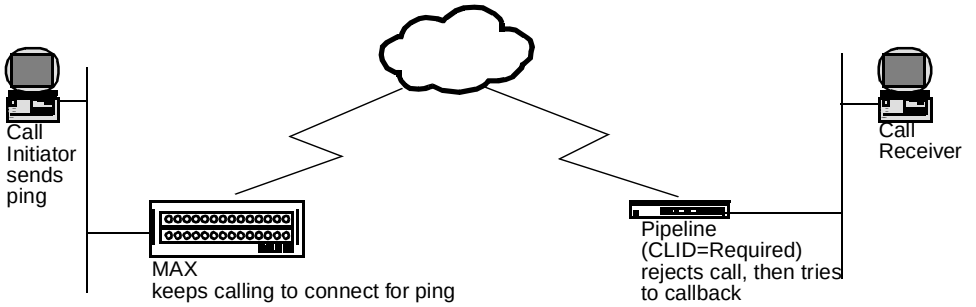

Figure 6-1. Callback connection failure

This is a special problem for ping and telnet, because these processes try continuously to open a connection and reject any callback because the process is already trying establish a connection.

When Expect Callback is set to Yes, calls that dialout and do not connect (for any reason) will be put on a list that disallows any further calls to that destination for 90 seconds. This gives the far end an opportunity to complete the callback.

Enabling Expect Callback

Expect Callback should only be set to Yes (TRUE) in dialout profiles, and not used for incoming calls.

To set Expect Callback to Yes:

1 Open Ethernet > Connections > *any profile* > Telco.

2 Set Exp Callback to Yes.

Note: If a call fails for any reason, regardless of whether or not the called machine requires CLID and is attempting a callback, the call initiator will still have to wait 90 seconds before attempting to call the same number again if Expect Callback is set to Yes.

SNMP write security disabled by default

A new parameter, R/W Comm Enable, whose default is No, disables set commands. Prior to this software release, the default behavior was to allow SNMP set commands.

Enabling SNMP write security

SNMP set commands enable you to load and save the Pipeline configuration using TFTP, and to make changes to the unit's configuration. With this software release, SNMP set commands are not permitted by default.

A new parameter in this feature, R/W Comm Enable, enables you to specify that SNMP set commands are enabled. To enable SNMP set commands:

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

```
90-B00 Mod Config
 SNMP options...
  Read Comm=public
 >R/W Comm Enable=No
  R/W Comm=N/A
```

2 Set R/W Comm Enable=Yes.

When R/WComm Enable=No, the R/W Comm parameter is N/A.

Note: To use a set command, you must know the read-write community string, even if R/W Comm Enable is set to Yes.

R/W Comm Enable

Description: Enables and disables the use of SNMP set commands.

Usage: Press Enter to select Yes or No.

- Yes enables the use of SNMP set commands. To use a set command, you must know the SNMP read-write community string specified in the R/W Comm parameter.
- No disables the use of set commands.

  No is the default.

Parameter Location: Ethernet > SNMP Options > Mod Config

SNMP request authentication added

This feature introduces SNMP request authentication, including reply protection. This implementation of SNMP request authentication is compatible with standard SNMPv1 practices, and affects the router's interpretation of SNMP messages that use it. Previously, the Pipeline did not provide authentication of SNMP requests.

You can use SNMP for security-related operations, such as altering the operational state of the router (rebooting, loading configurations, etc.) or firewall configurations. Since existing SNMPv1 is basically insecure, this feature adds authentication to verify that SNMP requests are only acted upon when they are known to be produced by an authorized system, and then only if they are known to be of recent origin.

Authentication Elements

This feature uses four elements to authenticate SNMP packets:

- secret authentication key
- data to be authenticated
- time-dependent state variables (for reply protection)
- MD5 hash value calculated over the key, data, and time.

The data, time, and hash values are transmitted with the packet. This allows the

management station and router to verify that the packet has been produced by an authorized system, and that the packet not been altered or significantly delayed in transmission.

The MD5 hash guarantees a high likelihood that only a system that knows the secret authentication key generated the packet, while the time variables guarantee a high likelihood that an attacker did not collect an authenticated packet and transmit it at a time of its own choosing, after a significant delay.

Community name string changes

This feature changes the internal structure of the write community name string in the router's configuration. The original SNMPv1 definition of the community string is that it is a string of octets that is compared to a similar string in the receiving SNMP entity. If the string in the packet exactly matches a community string in the receiving entity, then the packet is considered authentic. Currently, existing community names are simply ASCII strings with no internal structure. You can change these names as part of the router's configuration.

Currently the defaults are:

Ethernet > Mod Config > SNMP Options > Read comm=public

Ethernet > Mod Config > SNMP Options > R/W comm=write

The new structure for the community name separates the name from a key with the vertical bar character:

Ethernet > Mod Config > SNMP Options > R/W comm=write|secretkey

This causes the router to require SNMP SET REQUEST packets to be authenticated, using “secretkey” as the shared (but not transmitted) secret.

An authenticated community string contains the following structure:

The community name

- a zero-valued octet
- a router-defined “magic cookie” created by the router
- the router's current uptime (expressed in seconds)
- 16 bytes of MD5 hash calculated over:
  - secret key
  - magic cookie
  - uptime

– protocol data unit (PDU), which is the data in the packet

This structure is interpreted by the router when it receives an SNMP packet, and is produced by the router in response to a received SNMP request that contained an authenticated community name.

Cookie and uptime variables

The combination of cookie and uptime are expected to be a unique expression of real time for all routers, even routers that do not update calendar time, such as the Pipeline 50 and Pipeline 75.

Uptime is incremented once per second of router operation. This means that between reboots, for example, the uptime changes in a predictable way. The router compares the value of uptime in the packet with its own value, and calculates the difference. Packets are considered to have a freshly calculated hash value when the difference is less than 10 seconds, which indicates that the packet is less than 10 seconds old.

When the router is rebooted, a randomization algorithm produces a new cookie when the first authenticated SNMP request packet arrives after a router is rebooted.

The combination of a cookie value which changes on every reboot and an uptime value which changes between reboots enables the router to guarantee that the MD5 hash value is effectively unique for each unique authenticated packet, and that every authenticated SNMP packet must differ from every other SNMP packet.

The cookie and uptime are not encrypted or hidden in any way. Only the secret key is hidden from unauthorized users.

How SNMP authentication works

1 When the router receives an SNMP authentication request, it examines the community name in the packet, comparing it against its known community names.

– If the community name in the packet matches a non-authenticated community name in the router, then the router behaves normally, interpreting the packet, and generating a response packet with the same community name.

- If the community name in the packet matches neither a non-authenticated community name nor an authenticated community name in the router, then the packet is silently dropped.
- If the community name in the packet matches an authenticated community name in the router, then the router proceeds to authenticate the packet.

2 The router calculates an MD5 hash over its secret key, the magic cookie, uptime, and PDU. If the calculated hash doesn't match the hash value in the packet, then the router drops the packet without further notification.

3 The router compares the values for magic cookie and uptime in the packet against its values for cookie and uptime.

4 If the packet's cookie is not identical to the router's cookie or the packet's uptime differs from the router's uptime by more than 10 seconds., the router returns a response packet containing the router's current cookie, current uptime, no PDU data (other than the sequence number, Errstat, and ErrIndex), and a hash value calculated over the secret key, cookie, uptime, and truncated PDU.

ErrStat is 99, indicating a “reply error” and ErrIndex is zero.

This response packet contains the router's current cookie and uptime so that the originator of the request can integrate the new information in a subsequent packet. The originator of the request produces a new packet that is reformulated to contain the updated cookie and uptime, with a new (different) hash value. This new packet should pass authentication tests.

5 The router receives and compares the new packet containing the updated values and new hash value.

6 When a packet has been fully authenticated by the router, the router produces the appropriate SNMP response (such as updated var-binds in a get-response, etc.).

The router includes its cookie and an updated copy of its uptime in the packet, which allows the originator of the packet to re-synchronize its perception of the router's uptime with the router.

Configuring SNMP Authentication

To configure SNMP authentication, enter the read-write community name in the R/W_comm parameter of the SNMP Options submenu of the Ethernet profile. The read-write community name should have the format

```
name|secretkey
```

where:

- `name` is the name you want to assign to the read-write community name.
- `secretkey` is the alphanumeric key used for authentication.
- a vertical bar separates the `name` from the `secretkey`.

SNMP Get retrieves MPP session statistics

For the Pipeline 75 and 130 only. MPP session statistics appear in the Dyn Stat status window. Now, you can use SNMP get requests to query these values.The mppActiveStatsTable, added to the systemStatusGroup of the Ascend MIB (.1.3.6.1.4.1.529.12) to provide the objects necessary for an SNMP get request for session statistics.

Using a Get request to obtain MPP session statistics

The user interface changes are within the SNMP get requests. For example, if you use a simple SNMP "walk" utility to perform a walk request on the object identifier `.1.3.6.1.4.1.529.12.4` you will obtain sets of values that correspond to those that appear in the Dyn Stats window for a single MPP session on the LCD display. Each set of values is assigned an MpID.

Note: A `walk` utility is a form of `get next` request that begins with the zero index. Since the zero index does not exist (the index begins at 1), the utility returns the first available index, which would normally be 1, and continues returning indexes until there are no more available indexes.

Value sets returned by a get request

The values in the table below appear in each set returned by an SNMP `walk` or `get` request on the mppStatsMpID. For more information on these parameters,

see the Reference Guide that came with your documentation.

Value in mppStatsTable	Dyn Stats Parameter	Description
mppStatsRemoteName	*profile*	Connection Profile name set up in the Pipeline for this connection; shown at top of Dyn Stats window.
mppStatsQuality	Qual	Second line of Dyn Status window shows the quality of the link. Possible values are: Good, Fair, Marg, Poor, and N/A (link is not online)
mppStatsStartingTimeStamp	*time*	The amount of time the link has been active. When the link has been active for more than 96 hours, the duration is reported in days.
mppStatsBandwidth	*data rate*	Third line of Dyn Stats window shows the current data rate
mppStatsTotalChannels	*n* channels	The number of channels the data rate in mppStatsBandwidth represents.
mppStatsCLU	CLU *n*%	Current line utilization.
mppStatsALU	ALU *n*%	Average line utilization.

SNMP helps associate a call with a device

For the Pipeline 130 only. When a user calls the support desk about a problem with a connection, the help desk's management application can now use SNMP to isolate the device the user is logged into.

Overview

A new parameter, in the Line Profile associates up to three hunt groups with the

T1 line. A network management application can obtain this information using new SNMP variables and store a table of devices and the hunt group numbers in their WAN Line Profiles. When a user calls in with a problem, you can use this table to isolate the device(s) associated with the hunt group number the user called.

Configuring the hunt group numbers

1. Open the Net/T1 Line Profile for any line associated with a hunt group.
2. Enter the phone number for up to three hunt groups to be associated with calls logging into the line in the Hunt-n #

```
10-1** Factory
>Line 2...
 Hunt-1 #=
 Hunt-2 #=
 Hunt-3 #=
```

3. Save your changes

Hunt-*n*#

Description: Indicates a hunt group from 1 to 3 associated with the T1 line in a specific Line Profile. An SNMP manager can retrieve these numbers from Ascend devices and store them in a table that includes the devices from which information is retrieved and the hunt group numbers in their WAN Line Profiles.

Usage: Enter the phone number for the hunt group associated with current line in the Hunt-x # parameter.

Example: Hunt-1 #=847-4747

Dependencies: The numbers entered in the Hunt-n # parameters must be the same as the numbers that are assigned to T1 channels, creating the hunt group

Parameter Location: Net T1 Line Profile > Line Config

SNMP Enhancements

The sysConfigTftpCmd { ascend systemStatusGroup sysConfigTftp 1 } now has the following values that can be set:

- tsave (1), save the current configuration to a file. This saves only non-default parameter values.
- trestore (2), upload a valid configuration from a file via TFTP.
- tsave -a (3), save the current configuration to a file. This saves all parameter values, even those with default values.
- tsave -m (4), save the current configuration to a file, using the MIB OID instead of the VT100 interface names. This saves only non-default parameter values.
- tsave -am (5), save the current configuration to a file, using the MIB OID instead of the VT100 interface names. This saves all parameter values, even those with default values.

The values returned from the { mib-2 system sysObjectID } Pipeline identification OIDs are:

- { ascend products pipeline 5 } for Pipe50
- { ascend products pipeline 6 } for Pipe75
- { ascend products pipeline 7 } for Pipe130

Fixed interfaces appear first in SNMP IfTable

Fixed entities, such as hardware entities, now appear in the IfTable before software entities. Since IfNumber is referenced in many places in the MIB, this change helps to make other tables appear more consistent than previously. This change is compatible with previous MIB releases, since it is static information.

Administration

7

Overview

The following new features might affect the way you administer your unit:

Display unwanted dial-out packets

A new diagnostic option captures and displays packets that cause the Pipeline to dial out. You can then use the information to write data or call filters to prevent the packets from bringing up unwanted connections.

This enhancement adds the wan-data dial-out (wdDialout) option to the diagnostic monitor.

When packets are not captured

If a dial out is initiated for any of the following reasons, the wdDialout option does *not* capture a packet:

- Dial out caused by the Ctrl-D user command
- Dial out caused by callback security
- Dial out on nailed channels
- Dial out caused by NAT (Network Access Translation) acquiring an IP address
- Dial out initiated for IP over X.25, when the X.25 internet profile changes to active and there is data waiting for X.25 to bring up the connection
- Dial out caused by IGMP (Internet Group Management Protocol) multicast forwarding
- Dial out to acquire a DNS address during PPP negotiations
- Dial out in response to a DHCP Discover message
- Dial out caused by the Pipeline sending a DHCP packet for DHCP client processing
- Dial out caused in response to an APP (Ascend Password Protocol) Connect Request message

Turning on the diagnostic option

1. Enter the diagnostic mode by quickly typing:
   Esc [Esc =
2. At the ">" prompt, type:

   ```
   help ascend
   ```

you should see the wdDialout option listed. By default, the option is off.

3 To turn the option on, type:

```
wdDialout
WANDATA dialout display is ON
```

This is a toggle command. Typing it again turns the option off. See the next section for details on how packets are displayed in the diagnostic monitor.

4 To exit the diagnostic mode and return to the VT100 interface, type:

```
quit
```

Displaying packets

You can view wdDialout output in the diagnostic monitor. This section shows several examples.

Example 1

In the following example, the Pipeline unit's time and date have not been explicitly set, either by user command or SNTP server. So, the date and time in the captured packet is invalid. The phone number dialed on receipt of this packet is 92233002.:

```
 Date: 01/01/1990.          Time: 00:00:53
Cause an attempt to place call to 92233002
WD_DIALOUT_DISP: chunk 260126 type OLD-STYLE-PADDED.
: 42 octets @ 2C6950
  [0000]: ff ff ff ff ff ff 00 c0 7b 61 44 fe 08 06 00 01
  [0010]: 08 00 06 04 00 01 00 c0 7b 61 44 fe cc b2 d7 7b
  [0020]: 00 00 00 00 00 00 cc b2 d7 13

  [0000]: ff ff ff ff ff ff 00 80 c7 5b e9 5b 08 06 00 01
  [0010]: 08 00 06 04 00 01 00 80 c7 5b e9 5b cc b2 d7 13
  [0020]: 00 00 00 00 00 00 cc b2 d7 16 00 00 00 00 00 00
  [0030]: 00 00 00 00 00 00 00 00 00 00 00 00
```

The type OLD-STYLE-PADDED means that the packet has a 14-byte MAC (Ethernet) header + datagram (ARP request message). The packet contents pro-

vide the following information::

```
destination MAC address   ff:ff:ff:ff:ff:ff
source MAC address        00:c0:7b:61:44:fe   /* 123 */
arp packet type           08:06
arp_hrd                   00:01               /* Ethernet 1 */
arp_prot                  08:00               /* IP=0x800 */
arp_hlen                  06                  /* hlen = 6 */
arp_plen                  04                  /* plen = 4 */
arp_op                    00:01               /* arp ARP_REQ */
arp_sha                   00:c0:7b:61:44:fe   /* 123 */
arp_spa                   cc:b2:d7:7b         /* 123 */
arp_tha                   00:00:00:00:00:00
arp_tpa                   cc:b2:d7:13         /* 19 */
```

Example 2

In this example, the phone number dialed on receipt of this packet is 92233002. The type OLD-STYLE-PADDED means that the packet has a 14-byte MAC (Ethernet) header + datagram. This is a broadcast IP RWHO message. :

```
Date: 01/01/1990.  Time: 00:00:56
Cause an attempt to place call to 92233002
WD_DIALOUT_DISP: chunk 260126 type OLD-STYLE-PADDED.
: 198 octets @ 296810
  [0000]: ff ff ff ff ff ff 00 80 c7 5b e9 5b 08 00 45 00
  [0010]: 00 b8 0d c3 00 00 3f 11 24 fa cc b2 d7 13 cc b2
  [0020]: d7 ff 02 01 02 01 00 a4 e5 8a 01 01 00 00 32 46
  [0030]: 5e 26 00 00 00 00 63 6d 61 72 69 6e 65 72 00 00
  [0040]: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  [0050]: 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00 00
  [0060]: 00 00 32 46 4a e3 74 74 79 63 32 00 00 00 72 79
  [0070]: 75 00 00 00 00 00 32 46 4b 35 00 00 02 59 74 74
  [0080]: 79 63 33 00 00 00 72 79 75 00 00 00 00 00 32 46
  [0090]: 4b 39 00 00 00 3d 74 74 79 63 34 00 00 00 72 79
  [00a0]: 75 00 00 00 00 00 32 46 4b 3e 00 00 00 97 74 74
  [00b0]: 79 70 30 00 00 00 72 79 75 00 00 00 00 00 32 46
  [00c0]: 5e 00 00 00 00 01
```

The packet contents provide the following information:

```
destination MAC address   ff:ff:ff:ff:ff:ff
source MAC address        00:80:c7:5b:e9:5b
source IP adddress        cc:b2:d7:13       /* 204.178.215.19 */
destination IP address    cc:b2:d7:ff       /* 204.178.215.255
sub                                      network broadcast */
```

Example 3

In this example, the phone number dialed on receipt of this packet is 92233002. The type OLD-STYLE-PADDED means that the packet has a 14-byte MAC header + datagram. This is a unicast IP ICMP echo packet message.:

```
Date: 01/01/1990.     Time: 00:01:13
Cause an attempt to place call to 92233002
WD_DIALOUT_DISP: chunk 260126 type OLD-STYLE-PADDED.
: 98 octets @ 291EC8
  [0000]: 08 00 20 1f 5b ce 00 80 c7 5b e9 5b 08 00 45 00
  [0010]: 00 54 0e 09 00 00 ff 01 66 10 cc b2 d7 13 cc b2
  [0020]: d7 16 08 00 f5 1b bb 07 98 00 37 5e 46 32 3a 48
  [0030]: 0d 00 08 09 0a 0b 0c 0d 0e 0f 10 11 12 13 14 15
  [0040]: 16 17 18 19 1a 1b 1c 1d 1e 1f 20 21 22 23 24 25
  [0050]: 26 27 28 29 2a 2b 2c 2d 2e 2f 30 31 32 33 34 35
  [0060]: 36 37
```

The packet contents provide the following information:

```
destination MAC address   08:00:20:1f:5b:ce
source MAC address        00:80:c7:5b:e9:5b
source IP address         cc:b2:d7:13      /* 204.178.215.19 */
destination IP address    cc:b2:d7:ff      /* 204.178.215.22 */
```

Example 4

In this example, the phone number dialed on receipt of this packet is 917007337921. Note that there is no MAC header. This is an IPX packet: a Get

Nearest Server Request with service type File Server (0004).:

```
Date: 01/01/1990.        Time: 00:01:43
Cause an attempt to place call to 917007337921
WD_DIALOUT_DISP: chunk 261022 type IPX.
: 34 octets @ 2C6AA0
  [0000]: ff ff 00 22 00 11 00 00 00 00 ff ff ff ff ff ff
  [0010]: 04 52 00 00 00 00 00 a0 24 be d5 84 40 09 00 03
  [0020]: 00 04
```

The packet contents provide the following information:

```
chksum            ff:ff
packet len        00:22       /* 34 */
Transport Control 00          /* 0 */
packet type        11          /* 17 NetWare Core Protocol Packet */
dest network      00:00:00:00
dest Node         ff:ff:ff:ff:ff:ff
dest Socket       04:52        /* Service Advertising Protocol*/
source network    00:00:00:00:00
source Node       00:a0:24:be:d5:84  /*physical addr of src Node*/
Source Socket     40:09         /*4000h-7fffh Dynamic socket*/
Sap operation     00:03         /* Get Nearest Server Request */
Sap Service Type  0:04          /* File Server */
```

Example 5

In this example, the phone number dialed on receipt of this packet is 92233002. The type OLD-STYLE-PADDED means that the packet has a 14-byte MAC header + datagram. :

```
Date: 01/01/1990.  Time: 02:40:35
Cause an attempt to place call to 92233002
WD_DIALOUT_DISP: chunk 260126 type OLD-STYLE-PADDED.
: 60 octets @ 2AE950
  [0000]: 00 80 5f 74 93 d5 00 80 c7 2f 32 4c 00 2a ff ff
  [0010]: 00 29 00 11 30 6c 6b 00 00 00 00 00 00 01 04 51
  [0020]: 82 c1 b6 bf 00 80 c7 2f 32 4c 40 03 22 22 3f 03
  [0030]: 01 00 16 00 02 15 01 ff ff ff ff ff
```

The packet contents provide the following information::

```
destination MAC address   00:80:5f:74:93:d5
source MAC address        00:80:c7:2f:32:4c
chksum            ff:ff
packet len        00:29        /*41*/
packet type       11            /*17 NetWare Core Protocol Packet */
dest network      30:6c:6b:00
dest Node         00:00:00:00:00:01
dest Socket       04:51         /*NetWare Core Protocol (NCP Pkt)*/
source network    82:c1:b6:bf
source Node       00:80:c7:2f:32:4c  /* physical addr of src Node */
Source Socket     40:03             /*4000h-7fffh Dynamic socket*/
```

Configure call blocking on failed connections

For the Pipeline 50 and 75 only. You can now block additional retry attempts after a specified number of failed connection attempts have been made, and control the length of time call blocking is in effect.

Overview

When a connection fails, the Pipeline continues to try to complete the connection. This feature enables you to specify the number of unsuccessful retry attempts an Pipeline can make before blocking further attempts to make that connection. After the specified number of attempts have been made and failed, the blocking timer starts. The Pipeline continues to block further calls (discard packets) for a the period of time you specify.

Configuring call blocking

1. Open the Session options submenu of the Connection Profile.
2. Set the number of retry attempts that the Pipeline will allow to the Connection Profile.by enter the number in Block calls after=.
3. Specify the length of time during which the Pipeline will continue to block calls to the Connection Profile in Blocked duration=

Parameter reference

Two new variables have been added to the Session submenu of the Connection Profile.

Block calls after

Description: Specifies how many unsuccessful attempts the Pipeline will make before beginning to block calls (discard packets).

Usage: Enter the number of connection attempts permitted before the Pipeline blocks calls (discards packets) for the connection. The maximum number you can enter is 65535 (65535 attempts). The default is 0.

Parameter Location: Session Options submenu of the Connection Profile.

See Also: Blocked duration

Blocked duration

Description: Specifies the number of seconds the Pipeline will block calls (discard packets).

Usage: Enter the number of seconds for the Pipeline to block all calls made to the connection. When this period has elapsed, the unit will again allows calls to this connection.

Parameter Location: Session Options submenu of the Connection Profile.

See Also: Block calls after

Traceroute command added to terminal server

A Traceroute command has been added to the terminal server interface, similar to the existing terminal server Ping command. Traceroute is intended for use in network testing, measurement and management. It is useful for locating slow routers and in diagnosing IP routing problems. It is available on all platforms that offer a terminal server interface and IP routing and Telnet or Rlogin.

Note: The Traceroute command is available from the terminal server interface if outgoing Telnet or Rlogin is enabled, or if the user has Operations security.

The Internet is a large and complex aggregation of network hardware, connected together by gateways. Tracking the route packets follow or finding the gateway that's discarding your packets can be difficult. The Traceroute command utilizes the IP protocol "time to live" field and attempts to elicit an ICMP Time Exceeded response from each gateway along the path to some host.

The Traceroute command syntax is:

```
traceroute [ -n ] [ -v ] [ -m max_ttl ] [ -p port ] [ -q
nqueries ]
[ -w waittime ] host [ datasize ]
```

Note: The only mandatory parameter is the destination host name or IP number.

Options are:

-n	Prints hop addresses numerically rather than symbolically and numerically (this eliminates a nameserver 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	This mandatory parameter specifies 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).

Note: The -r and -s options (present in the UNIX version of Traceroute) are not supported.

The Traceroute command attempts to trace the route an IP packet would follow to some Internet host by launching UDP probe packets with a small TTL (time to live) and then listening for an ICMP "time exceeded" reply from a gateway. Probes start with a TTL of one and increase by one until we get an ICMP "port unreachable" message (which means we got to the host) or hit the maximum TTL.

Three probes are sent at each TTL setting and a line is printed showing the TTL, address of the gateway 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, a "*" is printed for that probe.

We don't want the destination host to process the UDP probe packets so the destination port is set to an unlikely value, such as 33434.

Possible annotations after the time field are as follows:

!H	Host reached.
!N	Network unreachable.
!P	Protocol unreachable.
!S	Source route failed. This should not occur and may indicate that there is a problem with the associated device.
!F	Fragmentation needed. This should not occur and may indicate that there is 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.

New tsave command option: -a

The tsave -a command option supplies a listing of all parameter settings. To use tsave -a, you need access to a UNIX host with a TFTP server. To produce the listing, use Telnet to access the Pipeline unit. Enter Ctrl-D to get to the DO menu and select D - diagnostics. At the terminal server prompt, enter the command using the syntax shown below:

```
tsave -a nnn.nnn.nnn.nnn file.name
```

Where:

-a	Lists all the menu items in the software for the unit.
nnn.nnn.nnn.nnn	Is the local IP address of a UNIX host with a TFTP server.
file.name	Is the name of an empty file you create first in the TFTP boot directory of the UNIX host. Be sure you have read/write access to the file. (If you run into problems, the reason usually has to do with lack of read/write access.) The output file is written to the TFTP boot directory of the UNIX host.

New tsave command option: -m

By default, the text configuration file you can create using the tsave command contains the VT-100 interface parameter names. A new option, -m, has been added to the tsave command to allow you to save the configuration file with the MIB field numbers instead of the parameter names.

Note: The files created by tsave and tsave -m can both be restored by trestore.

All Ascend products support this new option.

To use the tsave command, you must first enter diagnostic mode by quickly typing this four-character sequence

```
Esc [ Esc =
```

Then, to save the configuration of the Pipeline with the MIB field numbers instead of parameter names, enter this command line:

```
tsave -m <ipaddr> <filename>
```

Consider this example:

```
tsave -m 200.253.164.100 all
```

This command line saves the entire configuration of the Pipeline with an IP address of 200.253.164.100 to a file called “all”.

Values are saved in the format:

OOOO:MMMM.FFFF

where

- OOOO represents the Occurrence number (if > 0),
- MMMM represents MIB Type (if > 0),
- FFFF represents the MIB field number (if MMMM > 0).

Example

For example, the following text file results from performing a tsave on a sample system:

```
START=FILT=900=0
Name=IP Call
```

```
In filter 01...Valid=Yes
Out filter 01...Valid=Yes
Out filter 01...Generic...Forward=Yes
Out filter 01...Ip...Forward=Yes
END=FILT=900=0
```

If you perform a tsave -m command, this file is output:

```
START=FILT=900=0
54.1=IP Call
1:54.2,55.1=Yes
1:54.3,55.1=Yes
1:54.3,1:55.4,55.2=Yes
1:54.3,1:55.5,55.2=Yes
END=FILT=900=0
```

Consider this line:

```
1:54.3,1:55.5,55.2=Yes (Out filter
01...Ip...Forward=Yes)
```

[Out Filter] This is the 1st Occurrence of the Out Filter array; “Out filter 01..." belongs to the 54th MIB type; and it is the 3rd field in that MIB type. Thus the MIB tag generated is 1:54.3.

[IP] This is the first occurrence of an IP filter; the “IP...” belongs to the 55th MIB type; and it is the 5th field in that MIB. Thus the MIB tag generated is 1:55.5.

[Forward] There are not multiple occurrences of this field so the occurrence number is 0; “Forward” belongs to the 55th MIB type; and it is the 2nd field of that MIB. Thus the MIB tag generated here is 55.2.

All three MIB tags are now assembled (comma separated) into a single MIB tag: 1:55.3,1:55.5,55.2.

Larger executable load images enabled

A new system for loading larger system executables enables you to use tloadcode and TFTP from the diagnostic monitor. Previously the redundant system images

stored in the “top” and “bottom” halves of flash memory did not permit system loads larger than 448 KB.

Loading a fat system executable

Fat loads are loads whose compressed size exceeds 448 KB for the Pipeline. These system loads require special download procedures, which are described below.

Downloading a fat load

A fat system load can only be downloaded via tloadcode (TFTP) from the diagnostic monitor. Fat loads cannot be downloaded via the console port.

An older system image that is less than the maximum 448 KB will still load in the same manner as previously. These loads are referred to as “thin” loads.

If your unit currently is using a thin load system version that is not fat-load aware, you will first need to upgrade your current thin system to make it fat-load aware. This thin system should be backed up on your pc in case of fat-load failure. See “Loading a thin system that is fat-load aware” on page 7-15.

To load a fat-load aware system executable using TFTP:

1. From the Telnet interface, access the diagnostic monitor by typing these characters in rapid succession:

   `Esc [ Esc = (or Control "d",` then select `"D-diag")`

2. At the > prompt, type:

   `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.

   For example, the command:

   `tloadcode tftp-server ascend.bin`

   will load a software ascend.bin into flash from the machine named tftp-server. The current configuration is also saved to flash before new code is received, as a precaution.

3. One of the following messages appears:

   The following message is displayed at the default rate of 9600 bps if the load is thin:

```
UART initialized
thin load: inflate
..................................................
...
starting system...
```

The following message appears at the default rate of 9600 bps if the load is fat:

```
UART initialized
fat load: inflate
..................................................
....
starting system...
```

This completes code load if you have no errors.

Loading a thin system that is fat-load aware

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 control is transferred to the boot ROM's Xmodem serial download routine. To recover from this error and load the fat system, you must load a thin system that is fat-system aware. This thin load is required here because the boot ROM knows nothing about the new fat load format and only supports the traditional thin load. This thin load is probably not the system you will actually run, but it must be loaded first as a stepping stone toward downloading the desired fat system over the ethernet via tloadcode.

1 Invoke your Xmodem software to load the thin load through the console port.

2 Start the download of a thin load using the tloadcode command.

```
>>>> tloadcode:
```

The output of tloadcode has been modified slightly. When you download a traditional thin load, the following appears on the diagnostic monitor screen:

```
> tload yourmachinename /loads/newload.bin
saving config to flash

.....................................
.
loading code from nnn.nnn.nnn.nnn:nn
file /loads/newload.bin...
thin load:
.....................................................
.....................
```

Newload is the name of the binary you are attempting to load. The change is the addition of the line "thin load:" between the mention of the file name and the long series of dots.

3 After you have finished loading the fat aware thin load, reboot the unit.

4 Download the fat load using the tloadcode command.

When you download a fat load, the following appears on the diagnostic monitor screen:

```
> tload yourmachinename /loads/newload.bin
saving config to flash
.....................................
loading code from nnn.nnn.nnn.nnn:nnn
file /loads/newload.bin...
fat load part 1:
.....................................................
.....................
fat load part 2:
.............................................
```

Note the "fat load part *x*:" messages. They notify you when the first and second halves of the fat load are being loaded.

Note: In certain rare circumstances, a customer might possess a fat load from an engineering release written in an older format. tloadcode will automatically detect the obsolete format and refuse to load it, displaying a message similar to the following:

```
> tload squiddly /loads/mhpt1bri-moldy.bin
saving config to flash
.....................................
```

```
loading code from nnn.nnn.nnn.nnn:nnn
file /ascend/mb4/rtr/mhpt1bri/mhpt1bri-fatty.bin...
obsolete fat load format--discarding downloaded data...
-------------------------------------------------------
----------------------
```

Future unsupported loads

In the future, if you attempt to load a system that does not use the fat load format introduced by this feature, the load will be rejected if your current system does not support the new format.

```
> tload yourmachinename /loads/oldload-moldy.bin
saving config to flash

.....................................

loading code from 192.168.1.82:69
file /ascend/mb4/rtr/mhpt1bri/oldload-fatty.bin...
incompatible fat load format--discarding downloaded data
```

New Telnet password verification failure trap

This feature reports the IP address of the Telnet client whose login attempts failed. The address is included in the security violation message issued whenever the maximum number of Telnet login attempts to a Pipeline has been exceeded.

How the trap has been modified

To Telnet into a Pipeline, a user must supply the appropriate password, which is then verified. If the user cannot supply the correct password, an SNMP trap message is sent to all SNMP clients enabled for SNMP security messages.

The message includes the following information:

- The session number for the attempted Telnet session.
- The IP address of the host (the Pipeline).
- The associated IP address of the Telnet client that attempted the connection.

The format of the message is as follows:

```
mm.mmm.mmm.mmm  Enterprise Specific Trap (15) Uptime: xx:xx:xx
Name.iso.org.dod.internet.private.enterprises.ascend.sessionStatus Group.
IpAddress: ttt.ttt.ttt.ttt
sessionStatusTable.sessionStatusEntry.ssnStatusUserIPAddress%d
```

Where:

`mmm.mmm.mmm.mmm`	Host's IP address
`ttt.ttt.ttt.ttt`	Telnet client's IP address
`%d`	attempted Telnet session number

This trap message already existed in the listing of traps as an authentication failure (RFC-1215 trap-type 4). An authenticationFailure trap signifies that the Pipeline sending the trap is the addressee of a protocol message that is not properly authenticated. The only change to the trap is the addition of the IP address of the station that failed authentication.

Show system version command added

A show revision command has been added to the terminal server command line options for the show command.

The Show Revisions command

The show revision command displays system type and version information for the system currently running on the Pipeline, including:

- system name
- build name
- release number of the loaded software

For example, typing

```
show revision
```

at the command line prompt would display information similar to:

```
Pipeline system revision: mhpt1bip 4.6Bp10
```

HELP for Show command includes show revision

You can display a list of the options available for the show command. When you request help for the show command by typing

```
show ?
```

at the command line, the list of options that appears now includes the new Show system revision command:

```
show revision       Display system revision.
```

More information in fatal error log

The fatal error log now details the reason for a system reset and no longer describes a reset as a fatal error.

Reset descriptions

Previously, the fatal error log listed system resets as fatal errors, rather than resets, and did not give a reason for the reset. The fatal error log now lists a reset specifically as a reset and gives the reason for the reset.

Example: Reset from an NVRAM command

For example, if you use the diagnostic NVRAMCLEAR command to reset a unit, you would see something like the following:

```
OPERATOR RESET:  Index: 99   Revision: 4.6Bp10
        Date: 08/04/1996.       Time: 22:31:19
        NVRAMCLEAR Reset from unknown in security
profile 1.
```

```
OPERATOR RESET:  Index: 99   Revision: 4.6Be0
        Date: 08/04/1996.       Time: 22:32:23
        NVRAM was rebuilt
SYSTEM IS UP:  Index: 100   Revision: 4.6Be0
        Date: 08/04/1996.       Time: 22:33:00
```

Example: RESET from the diagnostics screen

If you use the diagnostic RESET command, you might see the following:

```
OPERATOR RESET:  Index: 99   Revision: 4.6Bp10
        Date: 08/04/1996.       Time: 22:32:23
       DEBUG Reset from unknown in security profile 1.
SYSTEM IS UP:  Index: 100   Revision: 4.6Be0
        Date: 08/04/1996.       Time: 22:33:00
```

Example: Reset from Sys Reset

If you select Sys Reset from the Sys Diag submenu of the System profile, you might see the following:

```
OPERATOR RESET:  Index: 99   Revision: 4.6Bp10
        Date: 08/04/1996.       Time: 22:32:23
        MENU Reset from unknown in security profile 1.
SYSTEM IS UP:  Index: 100   Revision: 4.6Be0
        Date: 08/04/1996.       Time: 22:33:00
```

Exceptions to how messages appear

If only one message is allowed in NVRAM, the "SYSTEM IS UP" message will not appear.

User-definable port for Syslog messages

To allow you more flexibility in controlling ports in a Syslog host, you can now specify the port at which a remote Syslog host listens for Syslog messages from an Pipeline. This feature enables you to run multiple copies of the syslog daemon on the Syslog host, with Pipelines sending syslog messages to different ports.

Syslog messages include warning, notice, and CDR (Call Data Reporting) records from the local system logs that are sent to the Syslog host. The Syslog host is the station to which the Pipeline sends system log messages, and the Log Port is the port on the Syslog host at which the host listens for these messages.

Previously, the Syslog host was always assumed to listen at a well-known port (port 514). You could not specify a different port.

Configuring the Log Port

1. Open the Ethernet > Mod Config menu.

```
90-C00 Mod Config
 Log...
  Syslog=Yes
  Log Host=206.65.212.205
  >Log Port=514
  Log Facility=Local0
```

2. Make sure that Syslog is enabled and a Log Host IP address is specified.
3. Select Log Port and type the port number at which you want the Syslog host to listen for messages from this Pipeline.

   The default port is port 514.
4. Close the Mod Config menu and save your changes.

Log Port

Description: Specifies the destination port on a syslog host where an Ascend unit's syslog messages will be received.

Syslog messages include warning, notice, and Call Data Reporting (CDR) records from the unit's local system logs.

Each Ascend unit can specify a different port, enabling the host to manage a number of units.

Usage: Select the Log Port parameter and enter a port number. The Log Port is the port on the Syslog host where the messages are received. The default is 514.

Dependencies: The Syslog parameter must be set to Yes. The Log Host parameter must contain the IP address of the station that will receive the syslog messages.

Parameter Location: Ethernet > Mod Config > Log

See Also: Syslog, Log Host

Terminal server and diagnostic functions

Two new options, Termsrv and Diagnostics, have been added to the Do menu. Previously, the functions supplied by these options could be not be accessed through the menu interface.

Accessing the new parameters

Two menu items, E=Termsrv and D=Diagnostics, have been added to the Do menu. The permissions set in a user's profile determine whether these options are available to that user.

To display the DO menu, press Ctrl-D:

```
Main Edit Menu
DO...
0=Esc
P=Password
C=Close TELNET
E=Termsrv
D=Diagnostics
```

Accessing the terminal server in other menus

The DO E option in the Main Edit DO menu has the same function as the Term Serv option in the Sys Diag menu.

Accessing the terminal server using keystrokes

You can use the following keystroke sequence (Escape key, left square bracket, Escape key, zero) to access the Terminal Server

```
<Esc> [ <Esc> 0
```

Set system clock using SNMP

An object has been added to the Ascend MIB to enable you to set the system clock using SNMP.

Shutdown PPP calls on authentication timeout

By default, if authentication fails on a PPP connection because of a bad password or an authentication server timeout, the Pipeline gracefully shuts down the PPP connection by sending an LCP-CLOSE request to the dial-up user. When Windows 95 receives the LCP-CLOSE during authentication, it assumes a rejected password, and displays a message telling the user that his or her password is invalid.

Disc on Auth Timeout

Description: Enables you to specify whether the Pipeline gracefully shuts down the PPP connection after an authentication timeout.

Usage: Press Enter to toggle between Yes and No.

- Yes specifies that the Pipeline does not shut down cleanly, but simply hangs up a PPP connection on an authentication timeout.
- No specifies that the Pipeline shuts down a call gracefully on a authentication timeout.

No is the default.

Dependencies: If PPP=No, Disc on Auth Timeout=N/A.

Parameter Location: Answer profile: Ethernet > Answer > PPP Options

See Also: PPP

TFTP checks compatibility of downloaded files

With this release, the Pipeline compares the software to be TFTP-uploaded to the currently loaded software. If the platform or network interface does not match, the Pipeline aborts the upload and displays information about why the abort occurred. The Pipeline will bypass this check if you use the TFTP command with the -f flag.

This feature protects you from unknowingly uploading software that is incompatible with your Pipeline. Previously, you were able to upload any software to any Pipeline. If you uploaded an incompatible software load, the upload would fail and revert to the previously-loaded software, but you received no indication of why the upload failed.

This check is initiated by the currently-loaded software. If your Pipeline is using a version of software with this feature and you attempt to load an older version of software that does not have this feature, the upload will be aborted because the older software has no platform identifiers that the currently-loaded software uses to validate compatibility. In this case, you'll need to use TFTP with the -f flag to have the Pipeline upload the older software without the compatibility check.

Examples

In the following example, a user attempts to use TFTP to upload a Pipeline 50 software load (b.p50) to a newer Pipeline 75 running b.v2p75:

1. From the VT100 interface, user accesses the diagnostics monitor.
2. User enters the following command:
   ```
   tloadcode tftpserver.ascend.com b.p50
   ```
3. The Pipeline 75 displays the following information to the screen:

```
saving config to flash
....................................
.
loading code from tftpserver.ascend.com
file /tftpboot/b.p50...
thin load:
This load appears to be for another platform.
This load appears not to support your network interface
Download aborted. Use 'tloadcode -f' to force.
```

The Pipeline 75has compared the uploading file, b.p50 to its currently-loaded file, b.v2p75. This informational messages indicate that the user attempted to load an incompatible platform and an incompatible network interface.

In the following example, a user attempts to use TFTP to upload an old version of software (without this feature) to a Pipeline 75 that uses this feature:

1 From the VT100 interface, user accesses the diagnostics monitor.

2 User enters the following command:

```
tloadcode tftpserver.ascend.com b.p50
```

3 The Pipeline 75 displays the following information to the screen:

```
saving config to flash
............
loading code from tftpserver.ascend.com
file /tftpboot/b.p50...
thin load:
This load has no platform identifier. Proceed with caution.
Download aborted. Use 'tloadcode -f' to force.
```

Here, the user decides to force the upload. The following messages are displayed:

1 User enters the following command

```
tloadcode -f tftpserver.ascend.com b.p50
```

2 The Pipeline 75 displays the following messages:

```
Download forced by user...
..................................................
...>
```

If you download an older version of software to a newer Pipeline 50 or 75 unit, the unit will become disabled and will have to be returned to Ascend for reconsideration.

Configure port for Syslog messages

To allow you more flexibility in controlling ports in a Syslog host, you can now specify the port at which a remote Syslog host listens for Syslog messages from an Ascend unit. This feature enables you to run multiple copies of the syslog daemon on the Syslog host, with Ascend units sending syslog messages to different ports.

Overview

You can now specify the port at which a remote host listens for syslog messages from an Ascend unit. Syslog messages include warning, notice, and CDR (Call Data Reporting) records from the local system logs that are sent to the Syslog host. The Syslog host is the station to which the Ascend unit sends system log messages, and the Log Port is the port on the Syslog host at which the host listens for these messages.

Previously, the Syslog host was always assumed to listen at a well-known port (port 514). You could not specify a different port.

Configuring the Log Port

1. Open the Ethernet > Mod Config menu.

```
90-C00 Mod Config
 Log...
  Syslog=Yes
  Log Host=206.65.212.205
  >Log Port=514
  Log Facility=Local0
```

2. Make sure that Syslog is enabled and a Log Host IP address is specified.
3. Select Log Port and type the port number at which you want the Syslog host to listen for messages from this Ascend unit.

   The default port is port 514.
4. Close the Mod Config menu and save your changes.

SNMP can detect concurrent sessions

You can now use SNMP to detect concurrent sessions with a single user.

Overview

A new table, sessionActiveTable, has been added to the Ascend MIB that enables you to detect concurrent sessions with a single user.The MAX must obtain and cache the ssnStatusCallReferenceNumfrom the RADIUS server to be retrieved by the SNMP get request.

Changes to the Ascend MIB

- **sessionActiveTable** OBJECT-TYPE sessionActiveGroup 3
  STATUS mandatory
  DESCRIPTION "A list of active session entries.
  This table is similar to sessionStatusTable with invalid entries screened out and indexed by:
  ssnActiveCallReferenceNum.
  ssnActiveCallReferenceNum tracks
  ssnStatusCallReferenceNum of
  sessionStatusTable."
- **sessionActiveEntry** OBJECT-TYPE sessionActiveTable 1
  SYNTAX SessionActiveEntry
  ACCESS not-accessible
  STATUS mandatory
  DESCRIPTION "An entry containing object variables to describe an active session."
- **ssnActiveCallReferenceNum** OBJECT-TYPE sessionActiveEntry 1
  SYNTAX INTEGER (1..'7fffffff'h)
  ACCESS read-only
  STATUS mandatory
  DESCRIPTION "A unique number identifying this active session.
  Refer to ssnStatusCallReferenceNum for more information."
- **ssnActiveIndex** OBJECT-TYPE sessionActiveEntry 2

SYNTAX INTEGER
ACCESS read-only
STATUS mandatory
DESCRIPTION "The index number for this session status entry. Its value ranges from 1 to 'ssnStatusMaximumSessions'. Refer to ssnStatusIndex for more information."

- **ssnActiveValidFlag** OBJECT-TYPE sessionActiveEntry 3
  SYNTAX INTEGER {
  invalid(1),
  valid(2)
  ACCESS read-only
  STATUS mandatory
  DESCRIPTION "All entries will be valid(2). Refer to ssnStatusValidFlag for more information."
- **ssnActiveUserName** OBJECT-TYPE sessionActiveEntry 4
  SYNTAX DisplayString
  ACCESS read-only
  STATUS mandatory
  DESCRIPTION "The name of the remote user. Refer to ssnStatusUserName for more information."
- **ssnActiveUserIPAddress** OBJECT-TYPE sessionActiveEntry 5
  SYNTAX IpAddress
  ACCESS read-only
  STATUS mandatory
  DESCRIPTION "The IP address of the remote user. Refer to ssnStatusUserIPAddress for more information."
- **ssnActiveUserSubnetMask** OBJECT-TYPE sessionActiveEntry 6
  SYNTAX IpAddress
  ACCESS read-only
  STATUS mandatory
  DESCRIPTION "The subnet mask of the remote user. Refer to ssnStatusUserSubnetMask for more information."
- **ssnActiveCurrentService** OBJECT-TYPE sessionActiveEntry 7
  SYNTAX INTEGER {

none(1),
other(2), -- none of the following
ppp(3), -- Point-To-Point Protocol
slip(4), -- Serial Line IP
mpp(5), -- Multichannel PPP
x25(6), -- X.25
combinet(7), -- Combinet
frameRelay(8), -- Frame Relay
euraw(9),
euui(10),
telnet(11), -- telnet
telnetBinary(12), -- binary telnet
rawTcp(13), -- raw TCP
terminalServer(14), -- terminal server
mp(15) -- Multilink PPP
ACCESS read-only
STATUS mandatory
DESCRIPTION "The current service provided to the remote user. The value none(1) is returned if entry is invalid OR if user dials into the terminal server and is in midst of a login sequence. Refer to ssnStatusCurrentService for more information."

SNMP can obtain active call status

A new table, the callActiveTable (1.3.6.1.4.1.529.11), has been added to the Ascend MIB. The new table enables you to use SNMP to obtain a listing of all active call-status entries. The information for each call in the listing corresponds to that in the Call Status window, described in the reference guide in your MAX documentation package.

To implement the new table, the following objects have been added to the Ascend MIB:

- **callActiveTable** OBJECT-TYPE callStatusGroup 16
  SYNTAX SEQUENCE OF CallActiveEntry
  ACCESS not-accessible

STATUS mandatory
DESCRIPTION "A list of active call status entries."

- **callActiveEntry** OBJECT-TYPE callActiveTable 1
  SYNTAX CallActiveEntry
  ACCESS not-accessible
  STATUS mandatory
  DESCRIPTION "An entry containing object variables to describe an active call's status."
- **callActiveCallReferenceNum** OBJECT-TYPE callActiveEntry 1
  SYNTAX INTEGER (1..'7fffffff'h)
  ACCESS read-only
  STATUS mandatory
  DESCRIPTION "The unique number identifying the session for which this call is associated."
- **callActiveIndex** OBJECT-TYPE callActiveEntry 2
  SYNTAX INTEGER
  ACCESS read-only
  STATUS mandatory
  DESCRIPTION "The index number for this call status entry. Its value ranges from 1 to 'callStatusMaximumEntries'."
- **callActiveValidFlag** OBJECT-TYPE callActiveEntry 3
  SYNTAX INTEGER {
  invalid(1),
  valid(2)
  ACCESS read-only
  STATUS mandatory
  DESCRIPTION "valid(2) for all active calls."
- **callActiveStartingTimeStamp** OBJECT-TYPE callActiveEntry 4
  SYNTAX INTEGER
  ACCESS read-only
  STATUS mandatory
  DESCRIPTION "The starting time for this call in seconds since startup."

- **callActiveDataRate** OBJECT-TYPE callActiveEntry 5
  SYNTAX INTEGER
  ACCESS read-only
  STATUS mandatory
  DESCRIPTION "The data rate for ISDN calls or the baud rate for modem calls."
- **callActiveSlotNumber** OBJECT-TYPE callActiveEntry 6
  SYNTAX INTEGER
  ACCESS read-only
  STATUS mandatory
  DESCRIPTION "Identifies the slot of the line being used. Its value ranges between 1 and the value 'slotNumber' in Ascend's slots group. This variable is equivalent to 'slotIndex' in the slot group."
- **callActiveSlotLineNumber** OBJECT-TYPE callActiveEntry 7
  SYNTAX INTEGER
  ACCESS read-only
  STATUS mandatory
  DESCRIPTION "Identifies the line for network slots. This variable is equivalent to 'slotItemIndex' in Ascend's slot group."
- **callActiveSlotChannelNumber** OBJECT-TYPE callActiveEntry 8
  SYNTAX INTEGER
  ACCESS read-only
  STATUS mandatory
  DESCRIPTION "Identifies the channel for the particular line identified by 'callActiveSlotLineNumber'."
- **callActiveModemSlotNumber** OBJECT-TYPE callActiveEntry 9
  SYNTAX INTEGER
  ACCESS read-only
  STATUS mandatory
  DESCRIPTION "Identifies the modem slot on the device. Its value ranges between 1 and the value 'slotNumber' in Ascend's slot group."
- **callActiveModemOnSlot** OBJECT-TYPE callActiveEntry 10
  SYNTAX INTEGER
  ACCESS read-only

STATUS mandatory

DESCRIPTION "Identifies the particular modem within a modem slot. A value of 0 indicates modems are not involved for this call."

- **callActiveIfIndex** OBJECT-TYPE callActiveEntry 11
  SYNTAX INTEGER
  ACCESS read-only
  STATUS mandatory
  DESCRIPTION "The interface index, ranging from 1 to the number of interfaces specified in the MIB-II variable ifNumber. The interface identified by a particular value of this index is the same interface as identified by the same value if ifIndex."
- **callActiveSessionIndex** OBJECT-TYPE callActiveEntry 12
  SYNTAX INTEGER
  ACCESS read-only
  STATUS mandatory
  DESCRIPTION "The index of the associated session entry. Value ranges from 1 to 'ssnActiveMaximumSessions'."
- **callActiveType** OBJECT-TYPE callActiveEntry 13
  SYNTAX INTEGER {
  callOutgoing(1), -- outgoing call
  callIncoming(2) -- incoming calL
  ACCESS read-only
  STATUS mandatory
  DESCRIPTION "Differenciates between outgoing and incoming calls."

Data rates reported to syslog

Two optional fields in the call-cleared Syslog posting show the transmit and receive data rates of the call. If the data rate is known, it is reported in bits per second after the “p” progress code, using the following identifiers:

- `s` (shows the call’s transmit rate)
- `r` (shows the call’s receive rate)

For example, this example output shows two messages reporting the data rates in

syslog.

```
...  ASCEND: call 1 AN slot 3 port 1 VOICE
...  ASCEND: call 2 AN slot 9 port 1 56KR
...  ASCEND: call 2 CL 0K  u=Torning c=185 p=60 s=64000
r=64000
...  ASCEND: slot 9 port 1, line 1, channel 1, Call Discon-
nected
...  ASCEND: call 1 CL 0K   c=20 p=40 s=31200 r=33600
...  ASCEND: call 3 AN slot 3 port 2 VOICE
...  ASCEND: call 3 CL 0K   c=185 p=31
...  ASCEND: slot 3 port 2, line 1, channel 2, Call Discon-
nected
```

If the data rate is not known it is omitted, as shown in the last three lines of the example Syslog output immediately above, where a call was placed to the Ascend device but no connection was made.The practice of omitting the data rate where not relevant is in accord with the handling of other fields in the same message.

Pipeline 75 Voice Features

A

Overview

The following new features might affect the way you use voice features on your unit:

Status display for voice calls

For the Pipeline 75 only. The 10-100 status window for the Pipeline or Pipeline 75 shows whether a voice call is on hold, as described in the following section.

Monitoring telephone connections

The status menu labeled 10-100 shows whether either or both of the B channels for your ISDN line is being used. An asterisk (`*`) to the right of B1 or B2 indicates that the channel used either for a voice or data call. The letter `h` indicates that a voice call is on hold. The letter `D` indicates that a call is being dialed.

In this example, B1—the first B channel—is in use.

```
10-100 1
 Link  D
 B1    *
 B2
```

In this example, B2—the second B channel—has one voice call on hold as well as an active voice call.

```
10-100 1
 Link  D
 B1
 B2    h *
```

WAN LED lit for voice calls

For the Pipeline 75 only. The WAN LED on the front of the Pipeline 75 is lit when the ISDN line is being used for a voice or data call.

Support 2-channel call on one SPID

For the Pipeline 75 only. This feature lets customers of AT&T 5ESS NI-1 reuse the same channel endpoint suffix (CES) for a 2-channel call (one voice and one data) on the same service provider identifier (SPID). The feature is not available for DMS-100 NI-1 customers.

The same CES can be reused to support 2-channel calls under the following circumstances (this information is for provisioning the line):

- The call is data on a specific CES.
- Phone Number Binding is TRUE and the new call attempt is voice over a specific CES currently assigned to a data call.

This feature is not available to DMS users, since the above calling attempts are rejected by DMS-100 NI-1.

Configuring for 2-channel, single SPID calls

There are no user interface changes to support this feature, but you must set the parameters in the Configure profile as follows:

Parameter	Required setting
My Num A	valid phone number (same as in standard NI-1 configuration with 2 SPID)
SPID 1 *or* SPID 2	Either SPID must have a valid SPID. This value will automatically be copied to the other SPID parameter.
Data Usage	Must be set to `A`
Phone 1 Usage or Phone 2 Usage	Both of these fields must be set to one of the following: • `A` • `None`
Phone Num Binding	Must be set to `Yes`. This will force the Pipeline to use the correct SPID.

Call conferencing

For the Pipeline 75 only. If your ISDN service includes the Call Conferencing

feature, you can use the Pipeline to establish conference calls. Conference calls allow more than two callers to converse at the same time. If the Call Conferencing feature is available from your telephone company, it allows either three-way conference calls (which include you and up to two other callers) or six-way conference calls (which include you and up to five other callers).

Using call conferencing

To establish a conference call, follow these steps:

1. Call a person to include in the conference, or have that person call you.
2. Put the call on hold by quickly pressing and releasing your telephone's switchhook (the button that is depressed when you hang up the telephone).
3. Call another person to include in the conference, or have that person call you.
4. Add anyone on hold to the call by quickly pressing and releasing your telephone's switchhook twice.
5. To add more callers to the conference call, repeat steps 2-4.

A caller normally leaves a conference call by hanging up. You can also drop the most recently added caller to a conference call by following this step:

Quickly press and release your telephone's switchhook twice.

Caller ID supported

The number of the calling party is included in the ISDN BRI data stream. You can capture the number, unless the caller has blocked it, by attaching a CallerID device to the POTS ports on the unit.

IDSL voice call support from Pipeline 75 or TA

For the Pipeline 75 only. Ascend's ISDN Digital Subscriber Line (IDSL) card now supports voice calls from a Pipeline 75 or any ISDN terminal adapter (TA) that supports en-bloc dialing. This allows you to make voice calls from an ISDN device through an IDSL line to a Pipeline. The Pipeline can then route that call to the voice network.

IDSL voice support from a Pipeline 75 or any other ISDN device requires that the ISDN device support Q.931 en-bloc dialing. A unit that supports en-bloc dialing reports the dialed numbers in the set up message it sends to the device it is connecting to. The Pipeline unit with the IDSL card installed can then use this information to route the call to the voice network.

Configuring an IDSL voice call

To configure the Pipeline:

1. Select System > Sys Config
2. Set Trunk Groups to Yes.
3. Exit and save the System profile.
4. From the main Edit menu, select the BRI/LT > Line Config > first Line profile.
5. Select the number of the line you want to configure.
6. Set the B1 and B2 Slot number parameters to the number of the slot you want to route incoming calls to.

   For example, to route calls from the first B-channel to slot 2 (a T1/PRI line) set B1 Slot to 2.
7. Exit and save first Line profile.
8. Continue configuring all the lines you want to provide voice service for.

To configure the Pipeline or similar device:

1. Make sure the device supports en-bloc dialing.

   For the Pipeline, in the Configure menu set the Switch Type to Japan.
2. When dialing out, make sure to prepend the number you want to dial with the trunk group.

   For example, to reach the number 555-5555, dial 2-555-5555 to get out the second T1/PRI line. If you omit the trunk group, the call is treated as any other call and is terminated at the Pipeline.
3. After entering the last digit of the phone number, enter an “end” character to indicate to the Pipeline that you have entered the entire phone number.

   On a Pipeline, enter a pound sign (#) after the phone number; other ISDN devices may have different end characters.

Switch Type

Description: Specifies the network switch type that provides ISDN BRI service to the Pipeline.

A network switch is the central office switch or PBX that terminates the ISDN BRI line at the MAX and connects the MAX to the circuit-switched WAN. The connection is a switched circuit consisting of one or more channels.

Usage: Press Enter to cycle through the choices. Your choices differ depending on the profile and enabled options.

You can select one of the switch types listed in the following table:.

Table A-1. Configure Profile switch types

Switch type	Explanation
AT&T/P-T-P	AT&T Point-to-Point is the default.
AT&T/Multi-P	ATT&T Multipoint.
NTI	Northern Telecommunications, Inc. Use this setting if your switch is DMS-100 Custom.
NI-1	National ISDN 1.
NI-2	National ISDN-2
IDSL	Identical to AT&T Point-to-Point, but has support for Q.931 en-bloc dialing.
U.K.	United Kingdom: ISDN-2 Hong Kong: HKT Switchline BRI Singapore: ST BRI Euro ISDN countries: Austria, Belgium, Denmark, Germany, Finland, Italy, Netherlands, Portugal, Spain, Sweden This is identical to NET 3.
SWISS	Switzerland: Swiss Net 2

Table A-1. Configure Profile switch types (Continued)

Switch type	Explanation
NET 3	This is identical to U.K.
GERMAN	Germany 1TR6 version: DBP Telecom
MP GERMAN	Germany: 1TR6 multipoint
FRANCE	France: FT Numeris
DUTCH	Netherlands 1TR6 version: PTT Netherlands BRI
BELGIUM	Belgium: Pre-Euro ISDN Belgacom Aline
JAPAN	Japan: NTT INS-64
AUSTRALIA	Australia and New Zealand

Dependencies: Keep this additional information in mind:

- The Switch Type parameter does not apply to a link using inband signaling (Call Type=56K or 56KR) or consisting entirely of nailed-up channels (Call Type=Nailed).

  For inband signaling, a line uses 8 kbps of each 64-kbps channel for WAN synchronization and signaling. The remaining 56 kbps handle the transmission of user data.

  Switched-56 lines use inband signaling.

- All international switch types except German operate in Point-to-Point mode.

Location: Configure Profile

Support for outgoing 3.1K audio calls added

For the Pipeline 75 only. Support for outgoing 3.1K audio calls has been added to allow calls to fax machines and other devices (particularly in Japan) that only accept a call indicated as a 3.1K audio call in the ISDN SETUP message.

How 3.1K audio calls work

In the Configuration menu, the current default is Phone 1 Usage=A and Phone 2 Usage=B. This means that the device type connected to both analog port 1 and analog port 2 are telephones. Any outgoing call from the corresponding analog port will use the Speech information transfer in its ISDN SETUP message.

If you set Phone 1 Usage=A 3.1K audio, the device type connected to analog port 1 is not a telephone and any call from the corresponding analog port will use the 3.1K audio information transfer in its ISDN SETUP message.

Note: Phone N Usage=A 3.1K audio and Phone N Usage=B 3.1K audio are for outgoing calls only. Many fax machines require Speech in the SETUP message, so the Pipeline will accept either Speech or 3.1K audio information transfer calls.

Configuring 3.1K audio call

To configure the Pipeline to send a 3.1K audio call:

1. Open the Configuration menu.
2. Set Phone N Usage=A 3.1K audio call.

   N is either 1 or 2, for phone 1 or phone 2. This specifies the analog port that will be used to make outgoing 3.1K audio calls.
3. Save the configuration.

Pipeline 130 Troubleshooting

B

Overview

The following new features might affect the way you troubleshoot your unit:

Backup Connection disconnect timer

For the Pipeline 130 only. When the nailed T1 connection on the Pipeline fails, the ISDN backup connection is established. This feature allows the ISDN backup connection to be timed out when the nailed T1 connection is reestablished.

The Pipeline provides a backup ISDN connection for its primary nailed T1 line. Previously, when the primary connection was reestablished after a failure, the backup ISDN connection had to be manually disconnected. This could create problems in fairly complex routing environments, because an ISP (for example) cannot completely prevent calls from their network over ISDN lines. Now the Pipeline senses when the primary connection has been reestablished and routes all traffic through this primary connection. This causes the ISDN line's idle timer to be activated, and the ISDN call is eventually terminated due to inactivity.

T1 loopback for the Pipeline 130

The Pipeline nailed T1 line now supports a loopback test. When the Pipeline T1 line is in loopback mode, it sends all the signals received from the switch back to the switch. Loopbacks can help diagnose whether the connection over the digital access line and the WAN is sound.

To support the T1 loopback test, these changes have been made in the Pipeline user interface:

- A new parameter, Loop Back, has been added to the Pipeline Mod Config > Nailed T1 menu.
- A new field has been added to the WAN status window

Descriptions of this parameter and status window field follow:

Loop Back

Description: Allows you to perform a loopback test of the Pipeline nailed T1 line. When the Pipeline T1 line is in loopback mode, it sends all the signals received from the switch back to the switch. Loopbacks can help diagnose whether the connection over the digital access line and the WAN is sound.

Usage: Press Enter to cycle through the choices.

- Normal specifies no loopback.
- Relay Loopback specifies a direct metallic loopback of the NI generated signal.
- Line loopback specifies a loopback of the network interface (NI) generated signal at the power level currently set in the Build out field of the Nailed T1 menu.

Dependencies: The T1 line cannot be used for communication when it is in loopback mode.

Parameter Location: Nailed T1 > Mod Config

New Line Status window field

The 10-100 Line Status window now contains a field below the T1/CSU field that indicates whether the T1 line is in loopback mode.

For example, if the Loop Back parameter in the Nailed T1 Mod Config profile is set to Normal, a Line Status window similar to the following is displayed

```
10-100 1   T1/CSU
Link   X   CARRIER
B1     *.....
B2     *.....
:
```

If Loop Back parameter in the Nailed T1 Mod Config profile is set to Loopback, a Line Status window similar to the following is displayed

```
10-100 1   T1/CSU
Link   X   LOOPBACK
B1     *.....
B2     *.....
:
```

Performing a loopback test

To perform a loopback test:

1 From the Main Edit menu select Nailed T1 > Mod Config.

2 Set the Loop Back parameter to Loopback.

3 Exit the Mod Config profile and save your changes.

4 As soon as you have saved the profile, the line is put into loopback mode.

Once the nailed T1 line is in loopback mode, no communication is possible over the WAN.

Manual loopback added for Pipeline 130

A manual loopback feature has been added to the Nailed 56 profile. This is a remote site loopback, where the signal is generated by the central office and looped back to the central office. This is different from Local loopback as used in POST, where the Pipeline signal loops back to itself.

The central office may request that the Pipeline user place the Pipeline in manual loopback mode when the central office is attempting to isolate problems. The default mode is Normal (not looped back), since no communication can occur while the Pipeline is in loopback mode.

Configuring manual loopback

To allow you to place the Pipeline in loopback mode, a field has been added to the Nailed 56 > Mod Config menu.

Note: Check the 10-100 Status Window to see if the Pipeline is in Normal or Loopback mode, as shown below, in "Line status window indicators for loopback mode."

To place the Pipeline in manual loopback mode:

1 Open the Nailed 56 profile.

If necessary press Escape until the main Edit menu is displayed, then select Nailed 56 and press Enter.

2 Select Mod Config.

The Mod Config menu appears, as shown.

```
30-100 Mod Config
 Nailed Grp=1
 Activation=Enabled
 Loop Back=Normal
```

3 Select Loop Back.

There are two possible values for this field: Normal or Loop—which is loopback mode.

4 Select Save, then press Enter to save the change.

Note: Remember to restore the Pipeline to Normal (not looped back) mode. Communication is not possible while in loopback mode.

Line status window indicators for loopback mode

The range of indicators in the 10-100 LCD status window has been increased to include a loopback mode indicator, as shown:

Status Indicator	Meaning
X	Line status down
A	Line status up
L	Line status looped back

Support for in-line loopback added for T1

Inband (in line) loopback is now functional for Pipeline 130 models with a T1 interface. Prior to this release, the Pipeline 130 ignored the inband loop command signal, therefore T1 technicians at the central office (CO) could not set a loopback on the line to diagnose problems.

Note: This loopback is the same one available by manual command in the T1 profile.

Changes in the WAN status window

The following changes can be seen in the WAN status window:

- The number of displayed line states has increased for the line state field.

- The definition of the loopback state has been expanded to include both manual and line loopback.

	Previous Display	New Display
RED	No Signal detected	No change
YELLOW	Yellow Alarm Signal detected.	No change
CARRIER	T1 CARRIER (LINE UP) detected.	No change
Loopback	Manual loopback active.	Manual loopback active (if you command it from the terminal) or line loopback active (if inband signal received from the network interface (sent by the phone company).
BLUE	(Previously detected, but not displayed.)	Blue Alarm Signal detected and displayed.

Manual T1 loopback using the line transceiver

Another loopback test has been added to the Pipeline 130 to perform manual full T1 loopback testing via the line transceiver.

Configuring manual loopback testing

1. Open the Mod Config > Nailed T1 menu.
2. Set Loop Back=Line Loopback

Traps for BRI linkUp and linkDown

Two BRI SNMP traps, which trigger alarm events or error events, are supported on the Pipeline 130. They are linkUp and linkDown. A trap is sent by the Pipeline 130 to indicate one of the following events has occurred:

- LinkUp indicates a BRI line has been physically connected to a BRI port while the Pipeline 130 is running, or a BRI line has been initialized during the boot/cold start process.
- LinkDown indicates a BRI line has been physically removed from a BRI port while the Pipeline 130 is running.

Index

A

B

C

D

E

F

G

H

I

L

M

N

P

Q

R

S

T

U

V

W