

MAX 6.0.0 Addendum

Ascend Communications, Inc.

Part Number: 7820-0149-002

For Software Version 6.0.0

March 11, 1998

Ascend is a registered trademark and Dynamic Bandwidth Allocation, MAX, MAX 200Plus, Multilink Protocol Plus, Pipeline, Secure Access Firewall, Global Digital Access are trademarks of Ascend Communications, Inc. Other trademarks and trade names in this publication belong to their respective owners.

Copyright © 1997–1998, Ascend Communications, Inc. All Rights Reserved.

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

Ascend Customer Service

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

Obtaining Technical Assistance

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

Information you will need

Before contacting Ascend Customer Service, gather the following information:

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

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

Telephone in the United States	800-ASCEND-4 (800-272-3634)
Telephone outside the United States	510-769-8027 (800-697-4772)
Austria/Germany/Switzerland	(+33) 492 96 5672
Benelux	(+33) 492 96 5674
France	(+33) 492 96 5673
Italy	(+33) 492 96 5676
Japan	(+81) 3 5325 7397
Middle East/Africa	(+33) 492 96 5679
Scandinavia	(+33) 492 96 5677
Spain/Portugal	(+33) 492 96 5675
UK	(+33) 492 96 5671
Email	support@ascend.com
Email (outside US)	EMEAsupport@ascend.com
Facsimile (fax)	510-814-2312
Customer Support BBS by modem	510-814-2302

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

Ascend Communications
1701 Harbor Bay Parkway
Alameda, CA 94502

Need information about new features and products?

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

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

`http://www.ascend.com`

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

`ftp.ascend.com`

Table of Contents

Ascend Customer Service.....	iii
Introduction.....	1
What is in this addendum	1
Related publications	1
A few words about this release.....	2
Upgrading system software	3
Known issues.....	18
WAN access features	19
BRI interface supports new analog encoding parameter	19
Support for 64K and 56K calls with R2 signaling.....	20
Israeli R2 signaling	20
Danish switch type added	21
MAX CH CNT upper limit specific to platform	22
B-Channel preference when MP+ or BACP adds bandwidth	22
BACP/BAP PPP protocol IDs match IETF.....	22
TACACS+ server retry attempts	23
Set Cause Code for ISDN DISCONNECT.....	23
ISDN Pre-T310 Timer	25
Enable data-service functionality for the MAX 4002 and MAX 4004 with AIM cards installed	
26	
Raw TCP connection enabled.....	26
TCP modem connections can be disabled	28
User-definable TCP connection retry timeout.....	29
IDSL voice call support.....	31
Loopback support for IDSL card (BRI/LT)	34
MAX supports new CSLIP Auto Detect parameter	37
More information provided for SLIP connections	38
Multilink or MP+ call now can span multiple MAX units.....	39
Updated Microsoft Callback Control Protocol support.....	46
MAXDial/IP support for immediate modem call restriction.....	50
Support added for multiple host selection	51
Limiting terminal server access per user	54
User-configurable call blocking after failed connection attempt	56
Specifying Shared Profiles per Connection Profile	57
Terminal server users can be forced to use unique profiles	58
CR-LF characters act as a single CR in terminal server.....	59
DNIS support for TCP-CLEAR, X25/PAD, and X25/T3POS calls.....	59
RADVision extensions to V.25bis for DNIS.....	60
Enhanced support for T1/PRI conversion	60
Enhanced support for outbound modem connections.....	64
PRI number plan in outgoing AIM/BONDING and PPP calls	65
PRI number plan in outgoing calls	67

IP routing features	67
Enable NAT on MAX 1800 and 2000	68
DHCP services on the MAX 1800 and 2000	82
Show DHCP command added to terminal-server	92
Changes to the IP router	95
Alphanumeric IP pool names	98
Removing routes to a host when the connection is down	99
Local DNS host address table option added	101
UDP Queue Control	107
Specifying the metric and preference for offline WAN connections.....	108
OSPF features	109
Specifying an OSPF ASE preference.....	109
Specifying the OSPF type of pool advertisement	110
Support for OSPF NSSAs (RFC 1587)	110
Advertise static routes as OSPF internal LSAs.....	112
Multicast features	113
Multicast default route	113
Configurable multicast group membership timeout.....	114
IPX features	114
NetWare SAP Home Server Proxy	115
SPX spoofing added for IPX.....	116
New Answer profile option for dial-in NetWare clients.....	116
Support for IPX without defining an IPX server	117
RADIUS features	118
ATMP attributes in RADIUS accounting	118
RADIUS refers to filter and firewall policies defined in local profile.....	121
Filter-Id (11).....	124
NAS-Port-Type added to RADIUS Accounting records	124
RADIUS Accounting checkpoint feature	125
NACK option for Local Profiles First parameter.....	126
Format of RADIUS NAS-Port attribute modified for the MAX	129
Call-logging feature	131
AppleTalk configurable from RADIUS.....	144
RADIUS Accounting ATMP Notification.....	147
Specifying preferences for routes configured in RADIUS	147
SecurID authentication for ARA (AppleTalk Remote Access) clients using RADIUS/ LOGOUT	149
MAX returns to primary RADIUS server after fixed time	149
New RADIUS attributes enable call routing to PPTP	151
RADIUS now recognizes packet types 33 and 34	153
RADIUS Accounting-Request packets show whether accounting is enabled or disabled	154
Configuring Data Over Voice Bearer Service (DOVBS) in RADIUS	155
New parameter for generating a second accounting Start record	157
New RADIUS and terminal server support for Point to Point Tunneling Protocol (PPTP)..	158
RADIUS support for IPX call and data filters	161

<i>Modem ID added to RADIUS Accounting Stop Records</i>	<i>164</i>
<i>RADIUS now follows NI-2 PRI spec for outgoing calls.....</i>	<i>165</i>
SNMP features.....	166
SNMP can enable/disable/quiesce T1/PRI links.....	166
Enable and Disable individual modem using SNMP.....	167
SNMP write security disabled by default	168
SNMP can help find which device a call is logged into	169
SNMP "get" now retrieves MPP session statistics.....	171
SNMP can monitor WAN lines and channels.....	173
SNMP can obtain active call status.....	178
SNMP system reset	180
SNMP can detect concurrent sessions	181
ASCEND MIB now includes modem status objects	183
SNMP RFC 1398 Ethernet-like MIB (dot3) support added	184
Set system clock through SNMP	184
Terminating user sessions using SNMP.....	185
Ascend MIB change supports counters for total and current calls	185
Firewall Control Protocol managed by SNMP	187
SNMP request authentication added	190
Initiating RADIUS updates using SNMP	192
SNMP sysConfigTftpStatus object reports more states	193
SNMP now reports reasons for last reset.....	193
Tunneling features	195
Layer 2 Tunnel Protocol (L2TP) support added	196
Maximum number of ATMP Tunnel sessions can be set	202
ATMP inactivity timer	203
Administration features.....	204
Support for Finger remote user information protocol (RFC 1288).....	204
Terminal server show users command added	205
Data rates reported to syslog	208
Receive and transmit rates now reported	209
Syslog enhancements	210
Syslog messages generated when Security profile activated.....	212
Configure port for Syslog messages	213
Defender authentication enhancements	214
RLOGIN enhanced -l option.....	215
TFTP checks for compatibility of downloaded files.....	216
Cexample now uses a session server key	217
T1-CSU Status window added to the MAX 2000.....	218
Modem features	218
Rockwell code version	218
New 56K digital modem cards.....	218
Disable modem capability added for MAX 1800 and MAX 2000	219
56k Modem Numbering	220
Ability to select between V.34 and K56Flex modulation.....	221

X.25 features.....	222
X.25 T3POS support	222
X.25 over the D channel	237
X.29 Reselection PAD message support.....	238
Answer Profile for X25/PAD terminal server users	238
Specifying X.3 parameters in an X.25 PAD Connection profile	239
X.25/IP inactivity timer supported.....	240
New X.25 user facilities parameters added.....	240
 Customized features.....	 241
Personal Handy Phone Service (PHS) support (Japan only)	241
T-Online	242
RADIUS bootup server supported for ZGR subaddresses and answer numbers.....	242
ZGR answer numbers obtained from RADIUS	246
ZGR subaddresses obtained from RADIUS	247
DTPT encapsulation for T-Online PPP sessions.....	248
PRI-PRI switching for T-Online.....	252
 Appletalk features.....	 254
AppleTalk routing added	254
Defender authentication added for AppleTalk Remote Access Protocol (ARAP)	263
Dial-in PPP support for AppleTalk	264
SecurID authentication for AppleTalk Remote Access (ARA) users.....	266
 Multiband features	 267
Enable data-service calls for the Multiband MAX 1800, 2000, and 4000.....	267
Suppress the display of the second T1 line on the Multiband MAX 2000	267
Suppress the T1-CSU and Serial WAN menus for the Multiband MAX 1800, 4002, and 4004 .	267
Suppress the T1-CSU and Serial WAN menus for the Multiband MAX 2000	268
Multiband simulation restricts MAX functionality.....	268
Multiband only option added	269

Introduction

This addendum applies to all MAX products except the MAX 200Plus unless otherwise noted in the text.

What is in this addendum

The documentation that came with your MAX unit describes how to install the hardware and configure the system. However, since the documentation was published, new system software has been released that contains features that are not yet included in the product documentation. This addendum describes those new features.

Related publications

Additional information is available in the MAX documentation set. The MAX documentation set consists of the following manuals:

- *MAX Getting Started Guide*. Explains how to install the MAX hardware. Includes the MAX technical specifications. A separate version of this guide is published for each MAX product and network interface (T1 or E1).
- *MAX ISP & Telecommuting Configuration Guide*. Explains how to use the VT100 interface to configure WAN connections and other related features. A separate version of this guide is published for each MAX product.
- *MAX Reference Guide*. An alphabetic reference to all MAX profiles, parameters, and commands. A separate version of this guide is published for each MAX product.
- *MAX Security Supplement*. Explains how to configure the MAX built-in security features. For information about configuring Secure Access Firewalls, see the documentation that came with your software. This guide is shared by all MAX products except the MAX 200Plus.
- *MAX RADIUS Configuration Guide*. Describes how to use RADIUS to configure WAN connections and other related features. This guide is shared by all MAX products except the MAX 200Plus.
- *MIF Supplement*. Explains how to use the Ascend Machine Interface Format (MIF), an alternative configuration interface for Ascend units. This guide is shared by all MAX products except the MAX 200Plus.

Introduction

A few words about this release...

A few words about this release...

Version 6.0.0 is the first product of a new software release process, which applies expanded internal testing designed to detect as many problems as possible prior to release. The new process adds layers of testing to Ascend's previous test suite.

We have endeavored to resolve any problems that might have major negative effects for your network, and to verify that the problems have been fixed. Some Trouble Reports (TRs) are still open, but this release resolves many of the highest priority issues. To fully complete our release process, we have added a formal beta test program, which you are invited to join for future releases, at www.ascend.com. This program and our newly expanded internal testing suite help ensure that each general release is tested as thoroughly as possible before we formally introduce it.

Our new software release process simplifies and unifies the version numbering we use for all products and integrates the MAX and Pipeline product lines into a single release.

Note: This release also introduces new software upgrade procedures, so please see the upgrade instructions in "Upgrading system software" on page 3.

A new section of the Release Notes, *Known issues*, describes open issues that might affect your environment. Some issues affect functionality. Others cause no functional problems, but can affect the VT100 display or terminal-server screen.

We know you'll see benefits from our new release process and renewed dedication to the quality of our releases.

Thank you,

Larry Gray
Director, Software Quality Assurance

Dana Harrison
Product Line Director, MAX products

Upgrading system software



Caution: The procedure for uploading new software to Ascend units have changed significantly. Carefully read the new software loading procedures explained in this section before upgrading your system.

This section explains how to upgrade your system software. It contains the following sections:

- Definitions and terms
- Guidelines for upgrading system software
- Before you begin
- Upgrading system software with a standard load
- Upgrading system software with a fat or thin load
- Upgrading system software with an extended load
- Upgrading system software from versions earlier than 4.6C to version 5.0A or above
- Using the serial port to upgrade to a standard or a thin load
- System messages

Definitions and terms

This document uses the following terms:

Build	<p>The name of the software binary.</p> <p>For example, <code>ti.m40</code> is the MAX 4000 T1 IP-only software build. For the names of all the software builds and the features they provide see</p> <p><code>/pub/Software-Releases/Max/Upgrade-Filenames.txt</code> or</p> <p><code>/pub/Software-Releases/Pipeline/Upgrade-Filenames.txt</code> on the Ascend FTP server.</p> <p>If possible, you should stay with the same build when upgrading. Loading a different build can cause your Ascend unit to lose its all or part of its configuration. If this happens, you must restore your configuration from a backup.</p>
Standard load	<p>Software versions 4.6Ci18 or earlier and all 4.6Cp releases. You can load these versions of software through the serial port or by using TFTP.</p> <p>TFTP is the recommended upgrade method for standard loads.</p>
Fat load	<p>4.6Ci19 to 5.0Aix and all 5.0Ap releases with a file size greater than 960 KB (for MAX units) or 448K (for Pipeline units). Before upgrading to a fat load for the first time, you must upgrade to a thin load.</p> <p>You must use TFTP to upgrade to fat loads.</p>

Thin load	4.6Ci19 to 5.0Aix and all 5.0Ap releases with a file size less than 960 KB (for MAX units) or 448 KB (for Pipeline units). TFTP is the recommended upgrade method for thin loads.
Restricted load	6.0.0 or later MAX release denoted by an “r” preceding the build name. For example, rti.m40 is the restricted load for the MAX 4000 T1 IP-only software build. Before upgrading to an extended load for the first time, you must upgrade to a restricted load. A restricted load only contains essential system software and is not meant to be run in a working environment. It does not have full functionality and is to be used only to upload to an extended load. TFTP is the recommended upgrade method for restricted loads. Pipeline releases do not have restricted loads.
Extended load	6.0.0 or later MAX release denoted by an “f” preceding the build name. You must use TFTP to upgrade to extended loads. For example, fti.m40 is the extended load for the MAX 4000 T1 IP-only software build. Pipeline releases do not have extended loads.

Guidelines for upgrading system software



Caution: Before upgrading, consider the following very important guidelines:

- Use TFTP to upgrade if possible. TFTP is more reliable and saves the Ascend unit configuration when you upgrade.
- You cannot load a fat load or an extended load through the serial port. You must use TFTP.
- If you are using TFTP to upgrade your software, use the `fsave` command immediately after executing the `tload` command. Failure to do so might cause your Ascend unit to lose its configuration.
- If possible, you should always stay with the same build of software when you upgrade. If you load a different version, your Ascend unit may lose its configuration. If this happens, you must restore your configuration from a backup.
- If you are upgrading to a software version 5.0A or 5.0Aix fat load for the first time, you must be on a load that supports the fat load format. All versions of software 5.0A or above support fat loads. You should perform the upgrade in two steps:
 - Upgrade to a thin load of the same build
 - Upgrade to the fat load
- If you are upgrading to a software version 6.0.0 or above, you must be on a load that supports the extended load format. All versions of software 6.0.0 or above support extended loads. You should perform the upgrade in two steps:
 - Upgrade to a restricted load of the same build
 - Upgrade to the extended load
- You can upgrade to a thin load or a restricted load from any version of software.

- If you are upgrading from software version 4.6C or earlier to software version 5.0A or later, see “Upgrading system software from versions earlier than 4.6C to version 5.0A or above” on page 12 for important information before you start.

Table 1 explains where to find the information you need to upgrade your unit.

Table 1. Ascend system software versions

Version you are upgrading to	Use the instructions in...
Standard load (4.6Ci18 or earlier and all 4.6Cp releases)	“Upgrading system software with a standard load” on page 6.
Fat or thin load (4.6Ci19 to 5.0Aix and all 5.0Ap releases)	“Upgrading system software with a fat or thin load” on page 7.
Extended load (6.0.0 or later)	A restricted load only contains essential system software and is not meant to be run in a working environment. It does not have full functionality and is to be used only to upload to an extended load. “Upgrading system software with an extended load” on page 10.

Before you begin

Make sure you perform all the tasks explained in Table 2 before upgrading your software.

Table 2. Before upgrading

Task	Description
If necessary, activate a Security Profile that allows for field upgrade.	If you are not sure how, see the section about Security Profiles in your documentation.
Record all of the passwords you want to retain, and save your Ascend unit’s current configuration to your computer’s hard disk.	For security reasons, passwords are not written to configuration files created through the serial console. A configuration file created using the Tsave command, however, <i>does</i> contain the system passwords. You can restore the Tsave configuration file using the serial console. If you chose to save your configuration using the serial console, you will have to restore your passwords manually. Restoring passwords is explained in “Using the serial port to upgrade to a standard or a thin load” on page 12.

Table 2. Before upgrading (continued)

Task	Description
Obtain the correct file, either by downloading it from the FTP server or by requesting it from Ascend technical support.	<p>To ensure that you load the correct software binary, you should check the load currently installed on your unit. To do so:</p> <ol style="list-style-type: none"> 1 Tab over to the 00-100 Sys Options window. 2 Press Enter to open the Sys Options menu. 3 Using the Down-Arrow key (or Ctrl- N), scroll down until you see a line similar to the following: Load: tb.m40 4 When upgrading, obtain the file with same name from the Ascend FTP site. <p>If your unit does not display the current load or you are unsure about which load to use, contact technical support.</p>
If you are upgrading to a fat load or an extended load for the first time, you must also obtain a thin load or a restricted load of the same build, if possible.	<p>For example, if you are upgrading a MAX 4000 to 5.0Ai13 fat load (such as tbim.m40), obtain a thin load of the same build (such as 5.0A tbim.m40).</p> <p>If you are upgrading to a MAX 6.0.0 extended load, obtain a 6.0.0 restricted load. Restricted loads are designated with an “r” in the load name. (For example rtbam.m40 is a restricted load).</p> <p>Newer Pipeline 50 or 75 units do not have fat loads and no Pipeline units have extended or restricted loads. Refer to /pub/Software-Releases/Pipeline/Upgrade-Filenames.txt to determine if you have a new Pipeline 50 or 75 unit.</p>
If you are using TFTP, make sure you load the correct binaries into the TFTP home directory on the TFTP server.	You must use TFTP to upgrade to a fat load or an extended load.
If you are using the serial port, make sure you have a reliable terminal emulation program, such as Procomm Plus.	<p>If you use the serial port, you can only upgrade to a standard or a thin load. Upgrading through the serial port is not recommended.</p> <p>If you use a Windows-based terminal emulator such as Windows Terminal or HyperTerminal, disable any screen savers or other programs or applications that could interrupt the file transfer. Failure to do so might cause the software upload to halt, and can render the Ascend unit unusable.</p>

Upgrading system software with a standard load

To upgrade system software with a standard load you can use either the serial port or TFTP. TFTP is the recommended method because it preserves your Ascend unit’s configuration. If you want to use the serial port to upgrade, see “Using the serial port to upgrade to a standard or a thin load” on page 12.

Using TFTP to upgrade to a standard load

To upgrade to a standard load using TFTP, you only have to enter a few commands. But you must enter them in the correct sequence, or you could lose the Ascend unit's configuration.

To upgrade to a standard load via TFTP:

- 1 Obtain the software version you want to upgrade to and place it in the TFTP server home directory.
- 2 From the Ascend unit's VT100 interface, access the diagnostics monitor by typing the following characters in rapid succession:

Esc [Esc =

Or, press Ctrl-D to invoke the DO menu and select D=Diagnostics.

- 3 At the > prompt, use the Tsave command to save your configuration as in the following example:

```
> tsave tftp-server router1.cfg
```

This saves the configuration of your unit to the file named `router1.cfg` in the TFTP home directory of the server named `tftp-server`. This file must already exist and be writable. Normally, TFTP upgrades save the configuration. Tsave is a precaution.



Caution: The file you save with the Tsave command contains all the passwords in clear text. You should move this file from the TFTP directory to a secure location after the upgrade procedure is complete.

- 4 Enter the following command:

```
tloadcode hostname filename
```

where *hostname* is the name or IP address of your TFTP server, and *filename* is the name of the system software on the server (relative to the TFTP home directory).

For example, the command:

```
tloadcode tftp-server t.m40
```

loads `t.m40` into flash from the machine named `tftp-server`.



Caution: You must use the Fsave command immediately after executing the Tload command. Failure to do so can cause your Ascend unit to lose its configuration.

- 5 Enter the following command to save your configuration to flash memory:

```
fsave
```

- 6 Enter the following command:

```
nvrampclear
```

After the Ascend unit clears NVRAM memory, it automatically resets.

This completes the upgrade.

Upgrading system software with a fat or thin load

Upgrading to a fat or thin load is not difficult, but you must be careful to follow the correct sequence of tasks.



Caution: If you are upgrading from software version 4.6C or earlier, see “Upgrading system software from versions earlier than 4.6C to version 5.0A or above” on page 12 for important information before upgrading.

To upgrade your system:

- 1 Obtain the software version binary you want to upgrade to and place it in the TFTP server home directory. If you are upgrading to a fat load for the first time, also obtain a thin load of the same build and place it in the same directory. (See page “Definitions and terms” on page 3 for an explanation of fat and thin loads.)



Caution: If possible, you should stay with the same build when upgrading. Loading a different build can cause your Ascend unit to lose all or part of its configuration. If this happens, you must restore your configuration from a backup.

For example, if you are upgrading a MAX 4000 to 5.0Ai13 fat load (such as `tbim.m40`), obtain a thin load of the same build (such as `5.0A tbim.m40`).

Note: Newer Pipeline 50 or 75 units do not have fat or thin loads, you only need to load a single software binary. Refer to `/pub/Software-Releases/Pipeline/Upgrade-Filenames.txt` on the Ascend FTP site to determine if you have a new Pipeline 50 or 75 unit.

- 2 From the Ascend unit’s VT100 interface, access the diagnostics monitor by typing the following characters in rapid succession:
Esc [Esc =
Or, press Ctrl-D to invoke the DO menu and select D=Diagnostics.
- 3 At the `>` prompt, use the `Tsave` command to save your configuration, as in the following example:

```
> tsave tftp-server router1.cfg
```

This saves the configuration of your unit to the file named `router1.cfg` in the TFTP home directory of the server named `tftp-server`. This file must already exist and be writable. Normally, TFTP upgrades save the configuration. `Tsave` is a precaution.



Caution: The file you save with the `Tsave` command contains all the passwords in clear text. You should move this file from the TFTP directory to a secure location after the upgrade procedure is complete.

- 4 At the `>` prompt, enter:

```
> tloadcode hostname filename
```

where *hostname* is the name or IP address of your TFTP server, and *filename* is the name of the system software on the server (relative to the TFTP home directory).



Caution: If you are upgrading from a standard load to a fat load, make sure you load a thin load first.

For example, the command:

```
> tloadcode tftp-server t.m40
```

loads `t.m40` into flash from the machine named `tftp-server`.



Caution: You must use the Fsave command immediately after executing the Tload command. Failure to do so may cause your Ascend unit to lose its configuration.

- 5 Enter the following command to save your configuration to flash memory:
fsave
- 6 Enter the following command:
nvrampclear
After the Ascend unit clears NVRAM memory, it automatically resets.
- 7 If you are upgrading to a thin load, you are done. If you are upgrading to a fat load, repeat the procedure, this time uploading the fat load binary.

After a successful upgrade, one of the following messages appears.

- If the load is thin:
UART initialized
thin load: inflate
.....
starting system...
- If the load is fat:
UART initialized
fat load: inflate
.....
starting system...

This completes the upgrade if you have no errors. If the upgrade is not successful, refer to "Recovering from a failed fat load upgrade" next.

Recovering from a failed fat load upgrade

If a fat load has a CRC (cyclic redundancy check) error, the following message appears:

```
UART initialized
fat load: bad CRC!!
forcing serial download at 57600 bps
please download a "thin" system...
```

Immediately after this message appears, the serial console speed is switched to 57600 bps, and the Ascend unit initiates an Xmodem serial download. To recover from this error and load the fat system, you must first load a thin system that is fat-load aware. Proceed as follows:

- 1 Activate your Xmodem software.
- 2 After you have finished loading the fat-aware thin load, reboot the unit.
- 3 Use the Tload command to download the fat load.
When you download a fat load, messages similar to the following appear on the diagnostics monitor screen:

```
> tload 192.168.1.82 tbam.m40
saving config to flash
.....
loading code from 192.168.1.82:69
file tbam.m40..
fat load part 1:
.....
```

.....
fat load part 2:
.....
The “fat load part *n*:” messages notify you when the first and second halves of the download begin.

Upgrading system software with an extended load

Your first upgrade to an extended load requires a preliminary procedure. You must first upgrade to a restricted load. A restricted load only contains essential system software and is not meant to be run in a working environment. It does not have full functionality and is to be used only to upload to an extended load. Note that Pipeline units do not have extended loads.



Warning: You cannot upgrade to extended loads using an IP over X.25 connection because restricted loads do not have X.25 support.



Caution: If you are upgrading from software version 4.6C or earlier, see “Upgrading system software from versions earlier than 4.6C to version 5.0A or above” on page 12 for important information before upgrading.

To upgrade your system:

- 1 Obtain the software-version binary you want to upgrade to and place it in the TFTP server home directory.

Extended loads are denoted by an “F” preceding the build filename.

- 2 If this is the first time you have upgraded to an extended load, obtain a restricted load of the same build and place it in the directory.

For example, if you are upgrading a MAX 4000 to an extended load (such as `tbam.m40`), obtain a MAX 4000 restricted load (such as `rtbam.m40`).

- 3 From the Ascend unit’s VT100 interface, access the diagnostics monitor by typing the following characters in rapid succession:

Esc [Esc =

Or, press Ctrl-D to invoke the DO menu, and select D=Diagnostics.

- 4 At the > prompt, use the Tsave command to save your configuration, as in the following example:

```
> tsave tftp-server router1.cfg
```

This saves the configuration of your unit to the file named `router1.cfg` in the TFTP home directory of the server named `tftp-server`. This file must already exist and be writable. Normally, TFTP upgrades save the configuration. Tsave is a precaution.



Caution: The file you save with the Tsave command contains all the passwords in clear text. You should move this file from the TFTP directory to a secure location after the upgrade procedure is complete.

- 5 At the > prompt, enter:

```
tloadcode hostname filename
```

where *hostname* is the name or IP address of your TFTP server, and *filename* is the name of the system software on the server (relative to the TFTP home directory).



Caution: If you want to upgrade your system for the first time to a software version 6.0.0 or later, you must first upgrade your system to a restricted load. Failure to do so can cause your Ascend unit to lose its configuration.

For example, the command:

```
tloadcode tftp-server rtbam.m40
```

loads the restricted load `rtbam.m40` into flash from the machine named `tftp-server`.



Caution: You must use the `Fsave` command immediately after executing the `Tload` command. Failure to do so can cause your Ascend unit to lose its configuration.

- 6 Enter the following command to save your configuration to flash memory:

```
fsave
```

- 7 Enter the following command:

```
nvrnclear
```

After the Ascend unit clears NVRAM memory, it automatically resets.

If you have downloaded the extended load, the upgrade is complete.

If you have loaded a restricted load, your system boots up in restricted mode. Restricted mode only allows you to load software. You cannot change or save profiles. While in restricted mode, the Edit menu displays the following banner:

```
* * RESTRICTED MODE * * *
```

If your system boots up in restricted mode, perform the following steps:

- 1 At the `>` prompt, enter:

```
tloadcode hostname filename
```

where *hostname* is the name or IP address of your TFTP server, and *filename* is the name of the extended load of system software on the server (relative to the TFTP home directory).

For example, the command:

```
tloadcode tftp-server ftbam.m40
```

loads the extended load `ftbam.m40` into flash from the machine named `tftp-server`.

- 2 Enter the following command:

```
nvrnclear
```

After the Ascend unit clears NVRAM memory, it automatically resets.

Your system will then boot up with the new version of software running.

Upgrading system software from versions earlier than 4.6C to version 5.0A or above

If you are upgrading from software version 4.6C or earlier to version 5.0A or later, perform the upgrade in the following order:

- 1 Load version 4.6Ci18, following the procedure in “Upgrading system software with a standard load” on page 6.
- 2 Load version 5.0A, following the procedure in “Upgrading system software with a fat or thin load” on page 7.
- 3 Load version 5.0Aix or 6.0.0, following the procedure in “Upgrading system software with a fat or thin load” on page 7 (for software versions 5.0Aix) or “Upgrading system software with an extended load” on page 10 (for software version 6.0.0).



Caution: Failure to follow this procedure might cause your Ascend unit to lose or corrupt its configuration, and could render the unit unusable.

Using the serial port to upgrade to a standard or a thin load



Caution: Uploading system software via the serial console overwrites all existing profiles. Save your current profiles settings to your hard disk before you begin upgrading system software. After the upgrade, restore your profiles from the backup file you created. Since the backup file is readable text, you can reenter the settings through the Ascend unit’s user interface. To avoid having existing profiles overwritten, use TFTP to upgrade your unit.



Caution: You cannot upload a fat load or an extended load using the serial port; it must be done using TFTP.

Upgrading through the serial port consists of the following general steps:

- Saving your configuration
- Uploading the software
- Restoring the configuration

Before you begin

Before upgrading your system through the serial port, make sure you have the following equipment and software:

- An IBM compatible PC or Macintosh with a serial port capable of connecting to the Ascend unit’s Console port.
- A straight-through serial cable.
- Data communications software for your PC or Mac with XModem CRC/1K support (for example, Procomm Plus, HyperTerminal for PCs or ZTerm for the Mac).



Caution: If you use a Windows-based terminal emulator such as Windows Terminal or HyperTerminal, disable any screen savers or other programs or applications that could

interrupt the file transfer. Failure to do so might cause the software upload to halt, and can render the Ascend unit unusable.

Saving your configuration

Before you start, verify that your terminal emulation program has a disk capture feature. Disk capture allows your emulator to capture to disk the ASCII characters it receives at its serial port. You should also verify that the data rate of your terminal emulation program is set to the same rate as the Term Rate parameter in the System Profile (Sys Config menu).

You can cancel the backup process at any time by pressing Ctrl-C.

To save the Pipeline configuration (except passwords) to disk:

- 1 Open the Sys Diag menu.
- 2 Select Save Config, and press Enter.
The following message appears:
Ready to download - type any key to start....
- 3 Turn on the Capture feature of your communications program, and supply a filename for the saved profiles. (Consult the documentation for your communications program if you have any questions about how to turn on the Capture feature.)
- 4 Press any key to start saving your configured profiles.
Rows of configuration information appear on the screen as the configuration file is downloaded to your hard disk. When the file has been saved, your communications program displays a message indicating the download is complete.
- 5 Turn off the Capture feature of your communications program.
- 6 Print a copy of your configured profiles for later reference.

You should examine the saved configuration file. Notice that some of the lines begin with START= and other lines begin with END=. A pair of these START/STOP lines and the block of data between them constitute a profile. If a parameter in a profile is set to its default value, it does not appear. In fact, you can have profiles with all parameters at their defaults, in which case the corresponding START/STOP blocks are empty. Make sure that there are no extra lines of text or characters either before START= or after END=. If there are, delete them. They could cause problems when you try to upload the file to the Ascend unit.

Uploading the software

To upload the software:

- 1 Type the following four-key sequence in rapid succession (press each key in the sequence shown, one after the other, as quickly as possible):
Esc [Esc -
(Press the escape key, the left bracket key, the escape key, and the minus key, in that order, in rapid succession.) The following string of Xmodem control characters appears:
CKCKCKCK
If you do not see these characters, you probably did not press the four-key sequence quickly enough. Try again. Most people use both hands and keep one finger on the escape key.
- 2 Use the Xmodem file-transfer protocol to send the system file to the Ascend unit.

Your communications program normally takes anywhere from 5 to 15 minutes to send the file to your Ascend unit. The time displayed on the screen does not represent real time. Do not worry if your communication program displays several “bad batch” messages. This is normal.

After the upload, the Ascend unit resets. Upon completion of the self-test, the Ascend unit’s initial menu appears in the Edit window with all parameters set to default values. This completes the upgrade.

If the upload fails during the transfer, try downloading another copy of the binary image from the Ascend FTP server and re-loading the code to the Ascend unit. If you still have problems, contact Ascend technical support for assistance.

Restoring the configuration

Under certain circumstances, the serial-port method might not completely restore your configuration. You should therefore verify that your configuration was properly restored every time you use this method. If you have many profiles and passwords, you should consider using TFTP to upgrade your software. (See “Using TFTP to upgrade to a standard load” on page 7.)

To restore the configuration, you must have administrative privileges that include Field Service (such as the Full Access Profile, for example). You use the Restore Cfg command to restore a full configuration that you saved by using the Save Cfg command, or to upload more specific configuration information obtained from Ascend (for example, a single filter stored in a special configuration file).

To load configuration information through the serial port

- 1 From the Ascend unit’s VT100 interface, access the diagnostics monitor by typing the following characters in rapid succession:
Esc [Esc =
Or, press Ctrl-D to invoke the DO menu, and select D=Diagnostics.
- 2 At the > prompt, enter the Fclear command:
> fclear
- 3 At the > prompt, enter the NVRAMClear command:
> nvramclear
This causes the system to reset. When it comes back up, proceed with restoring your configuration.
- 4 Enter **quit** to exit the Diagnostic interface.
- 5 Open the Sys Diag menu.
- 6 Select Restore Cfg, and press Enter.
The following message appears:
Waiting for upload data...
- 7 Use the Send ASCII File feature of the communications software to send the configuration file to the unit. (If you have any questions about how to send an ASCII file, consult the documentation for your communications program.)
When the restore has been completed, the following message appears:
Restore complete - type any key to return to menu
- 8 Press any key to return to the configuration menus.

- 9 Reset the Ascend unit, by selecting System > Sys Diag > Sys Reset and confirming the reset.

Restoring passwords

For security reasons, passwords are not written to configuration files created through the serial console. A configuration file created using the Tsave command, however, *does* contain the system passwords. You can restore the Tsave configuration file using the serial console.

After upgrading you may have to re-enter all the passwords on your system. If you edit your saved configuration file, however, and enter passwords in the appropriate fields (by replacing the word *SECURE* in each instance), these passwords will be restored. But note that if you do choose to edit your configuration file, you must save it as text only or you will not be able to load it into your unit.

If you restored a complete configuration, the passwords used in your Security profiles have been wiped out. To reset them:

- 1 Press Ctrl-D to invoke the DO menu, select Password, and choose the Full Access profile.
- 2 When you are prompted to enter the password, press Enter (the null password).
After you have restored your privileges by entering the null password, you should immediately open the Connection profiles, Security profiles, and Ethernet profile (Mod Config menu), and reset the passwords to their previous values.

System messages

Table 3 explains the messages that can appear during your upgrade.

Table 3. System software messages

Message	Explanation
UART initialized fat load: bad CRC!! forcing serial download at 57600 bps please download a "thin" system...	The fat load has a CRC (cyclic redundancy check) error. Immediately after this message appears, the serial console speed is switched to 57600 bps, and the Ascend unit initiates an Xmodem serial download. Load a thin load that understand the fat load format, as explained in “Upgrading system software with a fat or thin load” on page 7.
File tbam.m40 incompatible fat load format-- discarding downloaded data	You attempted to upgrade to a fat load from a version of system software that does not understand the fat load format. You must first load a thin load that is fat load aware, as explained in “Upgrading system software with a fat or thin load” on page 7.

Table 3. System software messages (continued)

Message	Explanation
This load has no platform identifier. Proceed with caution.	This message can occur if you are running software version 5.0Ai11 or later and you load an earlier incremental or patch release onto your system. The message indicates that Tloadcode cannot determine which platform the code is intended for. If you are using the correct software version, you can ignore this message.
This load appears not to support your network interface. Download aborted. Use 'tloadcode -f' to force.	Indicates you are attempting to load a version of code intended for a different network interface (for example, loading MAX 4000 T1 software onto a MAX 4000 E1 unit).
This load appears to be for another platform. Download aborted. Use 'tloadcode -f' to force.	Indicates you are attempting to load a version of code onto a platform for which it is not intended (for example, loading MAX 4000 software onto a MAX 2000). This is not recommended
UART initialized fat load: inflate starting system...	Indicates you have successfully loaded a fat load.
UART initialized extended load: inflate essential .+.+. invalid CRC!! entering restricted mode starting system...	Indicates the extended load has failed and that your system is being brought up in restricted mode. You must reload the software as explained in “Upgrading system software with an extended load” on page 10.
UART initialized extended load: inflate essential .+.+. invalid length!! entering restricted mode starting system...	Indicates the extended load has failed and that your system is being brought up in restricted mode. You must reload the software as explained in “Upgrading system software with an extended load” on page 10.
UART initialized extended load: inflate essential .+.+. inflate expendable..... starting system...	Indicates you have successfully loaded an extended load.

Table 3. System software messages (continued)

Message	Explanation
UART initialized thin load: inflate starting system...	Indicates you have successfully loaded a thin load.

Known issues

Known issues with this software release affect IPX, OSPF, modem out dial, and the display. There are also a couple of other issues, which will be resolved shortly.

IPX issues

Four known issues involve IPX. The first is that NetWare, by default, relies on the Data Link layer (also called Layer 2) to validate and guarantee data integrity. STAC link compression, if used, generates an eight-bit checksum, which is inadequate for NetWare data. This issue has existed in all previous releases of the software.

If your MAX supports NetWare (either routed or bridged), and you require link compression, you should configure your MAX in one of the following ways:

- Configure either STAC-9 or MS-STAC link compression, which use a more robust error-checking method, for any connection profile supporting IPX data. Configure link compression in the Ethernet > Answer > PPP Options > Link Comp parameter and Ethernet > Connections > *Any Connection profile* > Encaps Options > Link Comp parameter.
- Enable IPX-checksums on your NetWare servers and clients. (Both server and client must support IPX-checksums. If you enable checksums on your servers but your clients do not support checksums, they will fail to log in successfully.)
- Disable link compression completely by setting Ethernet > Answer > PPP Options > Link Comp = None and Ethernet > Connections > *Any Connection profile* > Encaps Options > Link Comp = None. By disabling link compression, the MAX validates and guarantees data integrity by means of PPP.

Following are the other IPX issues:

- When you bring up a call with the IPXPING command, *Answer* appears in the IPX Route Table as an entry under the *Origin* column. When the call is disconnected, the table returns to normal. This issue does not affect IPX functionality.
- The MAX does not send IPX RIP Response packets for a PPP-encapsulated dialup connection.
- If you support ATMP tunnels for your IPX clients, do *not* upgrade to 6.0.0.

OSPF issues

Under certain conditions we are currently unable to fully characterize, OSPF occasionally causes the MAX to reset. A fix for this problem will be available shortly in a maintenance release.

Modem out-dial issues

The following issues can affect modem out dial:

- In countries that support A-law encoding, 56K Modem handshake fails when a user dials out through either the MAX terminal server or the DeskDial client.
- When dialing a busy number, terminal-server out-dial modems do not report BUSY. The MAX detects the busy signal and goes on-hook before its modem receives the busy signal.

The MAX reports Cause Code 17 in the message log, the modem disconnects, and the screen displays NO CARRIER. This might be a problem for users or scripts that require a modem to report BUSY.

Display-only issues

The following issues cause display symptoms, but no functional problems:

- On MAX E1 units, channels 11 and 12 are missing their designators: 1 and 2, respectively. For example:

```
10-100 1234567890
L1/LA *****
3456789012345678901
***S*****_ _
```

- Occasionally, when you enter the terminal-server command Open, the MAX displays random characters on the screen.
- The MAX displays the Ethernet > Filters > NetWare Call filter for all loads, regardless of whether or not they support IPX.

Other issues

The following issues will be resolved shortly in a maintenance release:

- When Stacking is enabled, PAP-TOKEN-CHAP is not supported.
- Data filters are not supported when used in conjunction with Microsoft's CallBack Control Protocol (CBCP)
- Windows 95 machines installed with Winsock.dll (version 2.0) fail to log in successfully. You should configure Windows 95 in one of the following ways:
 - Upgrade Windows 95 with the latest Microsoft Dialup Networking patch 1.2, msdun12.exe.
 - Revert to older versions of Winsock.dll.

WAN access features

BRI interface supports new analog encoding parameter

Previously, the system chose the type of analog encoding for modem calls from the Switch Type parameter setting. This was the default. Now, you can override this Switch Type default by selecting Mu-Law or A-Law from the BRI Analog Encode parameter.

BRI Analog Encode

Description: Previously, the system chose the type of analog encoding for modem calls from the Switch Type parameter setting. This was the default. Now, you can override this Switch Type default by selecting Mu-Law or A-Law from the BRI Analog Encode parameter.

Usage: Specify one of the following values:

- SwitchType. This is the default. The system determines analog encoding for modem calls based on the Switch Type parameter setting.

You can override SwitchType by selecting either of the following two options:

- Mu-Law. The system uses mu-law analog encoding.
- A-Law. The system uses a-law analog encoding.

Example: BRI Analog Encode=Mu-Law

Location: Net/BRI > Line Config > BRI Analog Encode

Support for 64K and 56K calls with R2 signaling

The default bandwidth for data calls coming in over E1 channels using R2 signaling is now 64K. To configure a connection to use 56K instead, set the Force 56 parameter in the Telco Options subprofile of a Connection profile. When you set Force 56 to Yes, the MAX uses only the 56-kbps portion of a channel, even when all 64-kbps are available.

Following is an example of a procedure that specifies 56K for a call coming in over E1 channels using R2 signaling:

- 1 Open the Telco Options subprofile of a Connection profile.

```
Telco options...
>AnsOrig=Both
  Callback=No
  Exp Callback=No
  Call Type=Switched
  Group=N/A
  FT1 Caller=N/A
  Data Svc=64K
  Force 56=No
  Bill #=
  Call-by-Call=0
  Transit #=
  Dialout OK=No
```

- 2 Set Force 56 to Yes.

```
>Force 56=Yes
```

- 3 Close the Connection profile.

Israeli R2 signaling

R2 signaling is a CCITT standardized signaling protocol, which can be used on E1 digital trunks for establishing/clearing 64K switched circuits. R2 signaling is widely implemented in international markets where ISDN PRI is not yet available. The relevant specifications for

this signaling type are in ITU-T recommendations Q.400 to Q.490 and Israeli MFC-R2 Register Signaling documentation.

User interface changes

This section describes the changes to the user interface.

Grp B Answer Signal

Description: Specifies the group-B signal that the MAX sends immediately before answering an incoming call.

Usage: Specify B-1, B-2, and so on, up to B-15. The default is B-6, which is the recommended setting for E1-R2 Israeli signaling.

Systems in Mexico and Korea should set Grp B Answer Signal to B-1. Systems in Argentina should use B-6 (the default). For information about the proper settings for other countries, please contact your carrier.

Location: Net/E1>Line Config

See Also: Grp II Signal, Grp B Busy Signal

Grp B Busy Signal

Description: Specifies the group-B signal that the MAX sends as a busy signal.

Usage: Specify B-1, B-2, and so on, up to B-15. The default is B-3, which is the recommended setting for E1-R2 Israeli signaling.

Location: Net/E1>Line Config

See Also: Grp II Signal, Grp B Answer Signal

Danish switch type added

You can now specify a Danish network switch type for E1 units.

New parameter

Switch Type

Description: You can now specify a Danish network switch type in the Line Profile for MAX E1 units. The Danish switch type is a PRI ISDN switch and operates in point-to-point mode.

Usage: Select Switch Type=Danish.

Dependencies: You must set these additional settings for this parameter to work:

- Sig Mode=ISDN
- Framing Mode=G.703

Location: Net/E1 > Line Config > any profile

MAX CH CNT upper limit specific to platform

Now, the Pipeline or MAX will only allow you to configure Max Ch Cnt to a supported range of values. For example, the Max Ch Cnt parameter on a P50 can only be set to a maximum of 2. On the MAX 200+, it can be set to 8. On the MAX 4000, it can be set to 32.

Previously, the Max Ch Cnt parameter in each Connection Profile for every Pipeline or MAX unit was allowed any value from 0 to 32, regardless of whether the device supported 32 channels or not.

B-Channel preference when MP+ or BACP adds bandwidth

Previously, when an additional channel was requested due to an increase in traffic, the MAX randomly selected a B-channel for use. With this release, BACP and MP+ direct the MAX to give preference to the second B-channel of the BRI line on which the call originated.

Some Internet Service Providers offer discounts if multi-channel sessions use the same BRI line instead of spanning BRI lines.

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 4.

Previous Protocol ID	Current Protocol ID
8071	c02b.
0071	c02d

TACACS+ server retry attempts

TACACS+ backup server

This note adds the information to the documentation that the server makes two attempts to connect to each server.

How it works

The MAX sends a request for authentication to the first server on the list of hosts specified by Auth Host and waits for a response from the server for the number of seconds specified in the Auth Timeout parameter. If the MAX does not receive a response, within that time, it sends a second request for authentication to the same server and waits for the same amount of time. If the MAX does not receive a response within the specified timeout, it sends a request for authentication to the next server on the Auth Host list and repeats the process.

If the MAX is unsuccessful in obtaining a response from any of the servers on the list, the connection fails.

Set Cause Code for ISDN DISCONNECT

Overview

You can now specify that either User busy or Normal call clearing is sent as the Cause Element value in ISDN DISCONNECT packets when either Calling ID or Called ID Authentication fails.

There are two possible reasons the MAX sends out a DISCONNECT message with respect to ID authentication:

- due to a mismatch between the actual number and the expected number
- due to a RADIUS timeout

With this release, if either Calling ID or Caller ID authentication fails, the MAX can send out either of two Cause codes in the DISCONNECT message:

- User Busy
- Normal Call Clearing

WAN access features

Timeout Busy (previously CLID Timeout Busy)

The names and behaviors of two parameters in the Auth submenu of the Ethernet profile have been modified to support this feature. You can specify which Cause code the MAX sends out in DISCONNECT packets in these parameters:

- CLID Timeout Busy is now Timeout Busy
- CLID Fail Busy is now ID Fail Busy

Parameter Reference

Timeout Busy (previously CLID Timeout Busy)

Description: Specifies whether to return User Busy or Normal Call Clearing as a Cause in ISDN DISCONNECT messages when ID authentication fails due to a RADIUS timeout.

Usage: Press Enter to toggle between Yes and No. No is the default. If you choose Yes, and the ID authentication fails due to a RADIUS timeout, the DISCONNECT message will have the Cause value User Busy (decimal value 17). If you choose No, the Cause value will be Normal Call Clearing (decimal value 16).

Dependencies: This parameter will be N/A if Auth=None or Auth=TACACS+ in the this profile. The value set in this parameter applies to both Caller ID and Called ID authentication.

This parameter is N/A if ID Auth=Ignore.

Location: Ethernet Profile: Ethernet > Mod Config > Auth

See Also: IDFail Busy,

ID Fail Busy (previously CLID Fail Busy)

Description: Specifies whether to return User Busy or Normal Call Clearing as a Cause in ISDN DISCONNECT messages when authentication fails due to a mismatch between the actual number and the expected number.

Usage: Press Enter to toggle between Yes and No. No is the default. If you choose Yes, and the ID authentication fails due to a mismatch between the actual number and the expected number, the DISCONNECT message will have the Cause value User Busy (decimal value 17). If you choose No, the Cause value will be Normal Call Clearing (decimal value 16).

Dependencies: This parameter will be N/A if Auth=None or Auth=TACACS+ in the this profile. The value set in this parameter applies to both Caller ID and Called ID authentication.

This parameter is N/A if ID Auth=Ignore.

Location: Ethernet Profile: Ethernet > Mod Config > Auth

See Also: Timeout Busy,

Changes to Ascend MIB

Two values have been added to possible values for eventDisconnectReason in the Ascend MIB:

- **clidAuthFailed (value 4)**
This value indicates that the reason for the failure to authenticate is a mismatch between the expected value and the value provided by the caller.
- **clidAuthServTimeout (value 5)**
This value indicates that the reason for the failure to authenticate is a server timeout.

The Ascend MIB is updated frequently. You should always use the most current MIB. You can download the latest version of the Ascend MIB from the Ascend FTP server.

ISDN Pre-T310 Timer

This new feature allows the MAX to be configured to send either ISDN code 16 (Normal call clearing) or code 17 (User busy) when the T310 timer is triggered on PRI switch servicing the MAX. Previously, the MAX would return code 16 only.

Background

This feature is designed to allow callers into a MAX to get better clarification of call disconnects during the initial set-up of the call. If a call is presented to the MAX, and there is an extended period of delay while the call is being set up, for instance RADIUS requests or DNS lookups that are slowed due to a lot of local ethernet traffic, then you may want your users to get a different disconnect indication than the generic Normal call clearing.

In compliance with CCITT Specification Q.931, the MAX sends a CALL PROCEEDING message to the network switch for every call that is accepted..

The network switch sets its T310 timer as it awaits further messages from the MAX. The switch will tear down the call if the T310 timer expires. When this happens, the switch reports ISDN code 16 (Normal call clearing) to the calling device.

Previously, this was the only option. If the MAX cleared the call because there were no resources available to answer the call, there was no mechanism to inform the switch of the reason the call was not ultimately accepted.

This feature adds a MAX-specific timer which must be set to a time period less than the T310 timer on the switch. Then, after the MAX-specific timer expires but before the T310 timer expires, the MAX will send ISDN code 17 (User Busy) and clear the call.

Note: This feature is supported only on calls presented on T1/PRI lines.

Configuring the Pre-T310 timer

To configure the Pre-T310 timer:

- 1 Open the Net/T1 > Line Config > Line menu.
- 2 Set the Send Disc parameter to a value of from 0 to 60 seconds.
This must be set to a value less than the T310 timer value, so that it will expire (and the MAX will send its ISDN disconnect) before the T310 timer.
- 3 Open the Ethernet > Mod Config > Auth menu.

- 4 Set the Timeout Busy = Yes if you would like User Busy sent when the Send Disc timer expires. Set Timeout Busy = No if you would like Normal call clearing sent.

Note: This parameter was previously called CLID Timeout Busy.

Timeout Busy

Description: Determines the message sent by the MAX if a call clears due to the Send Disc timer expiring.

Usage: Press Enter to cycle through the choices (No is the default):

- Yes will cause ISDN code 17, User Busy, to be sent.
- No will cause ISDN code 16, Normal call clearing to be sent.

Send Disc

Description: This parameter specifies a number of seconds from 0 to 60. The value selected must be less than the T310 timer value used by the switch servicing the MAX.

Usage: Press Enter to open a text field. Then, type the number of seconds the MAX should wait from the time a call is presented to it before it clears the call. The timer is cancelled if the MAX sends a ISDN Alerting message or ISDN Disconnect message or if the network switch sends an ISDN Disconnect message. You can specify a number from 0 to 60. 0 disables this parameter. 0 is the default.

Enable data-service functionality for the MAX 4002 and MAX 4004 with AIM cards installed

For MAX 4002 and MAX 4004 units that have an AIM card installed, data-service functionality is now turned on. By default, the MAX 4002 and MAX 4004 have the data-service option turned off. In this release, if one of these platforms has an AIM card, the unit functions as though data service were turned on.

Raw TCP connection enabled

Users of the MAX's terminal server in the menu mode get a list of hosts to which they can telnet. With this software, that list can now configure include hosts to which the terminal server users can make a raw TCP connection.

Overview

In some instances, a terminal server user needs a raw TCP connection instead of a Telnet connection. This feature allows you to include the IP addresses and/or DNS names of a hosts to which the user can make a raw TCP connections.

Note: You cannot configure raw TCP hosts if you are using a RADIUS server to provide the list of hosts.

The IP address field for the TServ Options menu now holds a text string of up to 31 characters, instead of the dotted-decimal IP address format previously required.

Upgrading to a system software version containing this feature

To upgrade to a software version containing this feature, you must use nvramclear. This is required because the size of a field has changed. See the installation instructions that accompany new system downloads on the Ascend FTP server for more information on installing new system software.

Configuring a raw TCP host

- 1 Open the Ethernet > Mod Config > TServ options menu.
- 2 Select one of the Host # Addr fields and enter the following:

rawTcp hostname portnumber
rawTCP is the required string that causes the MAX to establish a raw TCP connection when the user chooses this host number. This entry is case-sensitive and must be entered exactly as shown.

hostname can be the DNS name of the host or the IP address of the host. The total number of characters, including the rawTcp string, must not exceed 31.

portnumber is the number of the port on which the connection for this host is to be established.
- 3 Enter a description of the host on the Host # Text field.

How it works

For example, assume the following entries in the TServ Options menu:

```
Remote Conf=No
Host #1 Addr=137.175.2.11
Host #1 Text=v
Host #2 Addr=
Host #2 Text=
Host #3 Addr=
Host #3 Text=
Host #4 Addr=rawTcp v 7
Host #4 Text=v 7
Immed Service=None
Immed Host=N/A
Immed Port=N/A
Telnet Host Auth=No
```

The Terminal Server menu would contain something like the following:

```
** Ascend Pipeline Terminal Server **

1. v
2. v 7

Enter Selection (1-2,q)
```

If a user picks host #4, a raw TCP connection is established to the host "v" on port 7.

WAN access features

Host #n Addr

If a user picks host #1, a Telnet connection is established to the host 137.175.2.11. In this case, the port is the default Telnet port.

Host #n Addr

Description: The Host #n Addr parameter has been modified to accept a string in the format of

rawTcp *hostaddress portnumber*

This strings indicate the IP address (or DNS name) of a raw TCP host and the UDP port that the TCP session to that host uses. As explained above the port number is optional.

TCP modem connections can be disabled

This feature allows you to enable or disable TCP modem access to the MAX as well as configure the default port for TCP modem access.

The MAX treats a TCP-encapsulated call between two MAX units over an asynchronous line as if it were a modem. This is referred to as TCP modem. Previously, the MAX would always allow such calls. Now, you can disable TCP modem connections to the MAX. In addition, you can change the TCP port used for these connections. Previously the default port for TCP modem access was 150. It has now been changed to 6150.

Figure 5 illustrates an example TCP modem setup. A user dialing into an ISP first connects to telephone switch, which then establishes a connection to a MAX. This local MAX has a TCP-Clear connection configured in RADIUS to a MAX at an ISP. Typically, this connection is over Frame Relay. The remote user appears to be directly connected to the ISP MAX; the local MAX merely passes the data through. All authentication of remote users is typically done by the ISP MAX.




Figure 5. Sample TCP modem connection

Parameter reference

This section describes the new MAX parameters.

TCP Modem Enabled

Description: Specifies whether the MAX allows TCP modem access.

Usage: Specify one of the following values:

- Yes indicates the MAX answers TCP modem connections.
- No indicates the MAX does not answer TCP modem connections over the port specified by TCP Modem Port.
No is the default.

Location: Ethernet > Mod Config > TCP Modem Options

TCP Modem Port

Description: Specifies the port for TCP modem access.

Usage: Specify a TCP port. The default is 6150.

Location: Ethernet > Mod Config > TCP Modem Options

User-definable TCP connection retry timeout

This feature enables you to set the maximum length of time a MAX tries to complete a connection with a host IP address before proceeding to the next address on the list provided by the DNS server. Previously, this timeout length was always 170 seconds, which is longer than some client software will permit before the client software times out.

When a terminal server attempts to connect to a host through a MAX, and a DNS server is used, the DNS server may supply a list of host IP addresses in response to the query from the MAX. The DNS server simply provides the list and has no way of determining whether the hosts at the addresses are available or how long it may take for them to respond.

If the DNS List Attempt feature is enabled, the MAX will attempt to connect to the first address on the list until the timeout expires. Previously, the timeout default value was 170 seconds, which means that the MAX would attempt to connect to the first address on the DNS list for that length of time. Some client software for the terminal server user times out before 170 seconds elapse. When the client software times out, the connection is dropped by the client and none of the remaining addresses on the DNS list are tried. The connection may never be successful because each time it is retried and the DNS server provides the same list of addresses, the MAX starts at the top of the list again and attempts to make the same connection that was previously unsuccessful.

A new parameter, **TCP Timeout**, has been added to the Ethernet Configuration Profile, shown in Figure 6. Set this parameter to a value between 1 and 200 to specify the TCP retry time so that connections to additional host addresses can be attempted, if necessary, before the client software times out. When the timeout value expires and the connection to an address is unsuccessful, the MAX will proceed to try the next IP address on the list for the length of time specified in TCP timeout.

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 MAX 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 MAX operate as previously and attempt for a maximum of 170 seconds to connect to each address on the list, until a connection is successful or the connection is dropped.

Figure 6. TCP timeout parameter in Ethernet Mod Config menu

```
90-A00 Mod Config
RIP Summary=Yes
ICMP Redirects=Accept
BOOTP Relay...
DNS...
Multicast...
Auth...
Accounting...
RADIUS server...
Log...
PPTP Options...
Modem Ringback=Yes
SNTP Server...
Stack Options...
UDP Cksum=No
TCP Timeout=0
Adv Dialout Routes=Always
```

TCP timeout parameter reference

Description: This parameter specifies the length of time during which a MAX will attempt to connect to an IP host in the list provided by the DNS server.

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

This feature applies to all TCP connections initiated from the MAX, including telnet, rlogin, tcp-clear, and the TCP portion of DNS queries.

Usage: To set the timeout value, select TCP timeout and type the number of seconds the MAX can attempt a connection to an IP address on the DNS list.

The range of values for TCP timeout 0 to 200 seconds. This specifies the number of seconds after which the MAX will stop attempting to connect to an IP address and will proceed to the next address on the list. (but as noted below, other limits already in the system may terminate

TCP retries after about 170 seconds). The number of start-connection messages the MAX will send is fixed, however.

Note: When the MAX 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 DNS Timeout has not yet elapsed.

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

The List Attempt parameter in the DNS submenu of the Mod Config menu in the Ethernet Profile must be enabled. This permits the MAX to attempt the IP addresses. On a list, if the DNS server provides such a list. The List Attempt parameter does not apply if Telnet and Immediate Telnet are both disabled.

Ethernet Profile: Ethernet/Mod Config

IDSL voice call support

With this release, Ascend's ISDN Digital Subscriber Line (IDSL) card (displayed as BRI/LT in the vt100 menu screen) supports incoming and outgoing voice calls. To support outgoing voice calls, the connected TE (Terminal Equipment) must send digits to the MAX using Q.931 en-bloc dialing (sends all dialed digits to the MAX in one block (the ISDN Call Setup message) rather than one digit at a time).

The MAX receives outgoing call requests from attached ISDN TE and routes voice calls to the PSTN (Public Switched Telephone Network) over a T1 line or ISDN PRI line. The MAX receives incoming voice calls and routes them to TEs connected to IDSL cards based on DNIS (Dialed Number Identification Service).

Configuring the MAX IDSL card for outgoing voice calls

To configure the MAX to accept voice calls from ISDN TEs connected to the IDSL slot card and route them to the PSTN network:

- 1 Open the System > Sys Config menu.
- 2 Set Use Trunk Groups to Yes.
- 3 Exit and save the System profile.

Use the following steps if you want voice call requests routed to a T1/PRI line:

- 1 Open the Net/T1 > Line Config > Line *n* menu.
- 2 Set Ch *n* TrnkGrp to a value from 4 to 9.

where *n* specifies the channel of the T1/PRI line you want to make available to the IDSL card.

You must prepend this value to the phone number the TE dials. When the MAX receives a voice call request from the TE, the MAX will use the trunk group number to route the call to a T1 channel with a matching trunk group number. If trunk groups are not used, the call request will terminate at the MAX and not be forwarded to the PSTN.

- 3 Exit and save Line profile.

For details on configuring your T1 line or PRI line, see the *MAX ISP and Telecommuting Configuration Guide*.

Configuring the MAX ISDL card for incoming voice calls

To configure the MAX to accept voice calls from the PSTN network and route them to TEs connected to the ISDL slot cards, select one of the following methods

The following instructs the MAX to route calls to the ISDL card on the basis of the called number:

- 1 Open the BRI/LT > Line Config > Line *n* menu.
- 2 Set Ans 1#, Ans 2#, or both to the called number that is dialed to reach the end user's TE. The Central Office (CO) switch must support DNIS since the MAX matches the DNIS number of the incoming call to configured numbers in Ans *n*#.

The following instructs the MAX to route calls to the ISDL card on the basis of the T1 channel on which the calls are received:

- 1 Open the Net/T1 > Line Config > Line *n* menu.
- 2 If a MAX should route calls received on a specific channel to the ISDL card, set the appropriate Ch *n* Slot parameter to the ISDL card's slot number.
For example, if the MAX is to route all calls received on channel 1 to an ISDL card in slot 7, set Ch 1 Slot to 7.

Note: For details on configuring your ISDL line and incoming call routing, see the *MAX ISP and Telecommuting Configuration Guide*.

Configuring the Pipeline for outgoing voice calls over ISDL

Use the following steps to configure the Pipeline to support outgoing voice calls when connected to a MAX ISDL slot card for routing to PSTN network:

Note: If you use a TE other than a Pipeline, it must support en-bloc dialing.

- 1 Open Ethernet > Answer > PPP Options menu
- 2 Set Encaps to MPP
MPP supports data call preemption. See Note: below.
- 3 Open the Configure menu.
- 4 Set Switch Type to ISDL.
The ISDL selection is an AT&T 5ESS Point-to-Point configuration with en-bloc dialing support.

When you dial out from a phone connected to the analog port of the Pipeline or TE, you must prepend the Trunk group number (configured on the MAX) to the phone number you dial. This is similar to dialing from an ISDN Centrex System, where you are required to prepend the phone number you dial with an additional digit to get an outside line.

For example, if the MAX is configured with Trunk Group set to 9 and you are dialing 555-5555, dial

9-555-5555 to instruct the MAX to dial 555-5555 on the channels (T1 or PRI) that are configured with a Trunk Group set to 9.

If you omit the trunk group, the call is terminated at the MAX. It is not routed to the PSTN.

Note: Data call preemption is supported with this feature. If you use two channels for a single MPP data call, and dial your analog phone, one channel will be reallocated to the voice call, leaving one channel for the data call. When you hang up, the channel will be reallocated to the data call if throughput load warrants it.

For additional information on call preemption, see the *Pipeline User's Guide*.

Modified parameter in the Pipeline

The Switch Type parameter supports the new value of IDSL.

Switch Type

Description: Specifies the network switch type that provides ISDN BRI service to the MAX.

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:.

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 Multitpoint.
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.

WAN access features

Loopback support for IDSL card (BRI/LT)

Configure Profile switch types (continued)

Switch type	Explanation
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
NET 3	This is identical to U.K.
GERMAN	Germany ITR6 version: DBP Telecom
MP GERMAN	Germany: ITR6 multipoint
FRANC	France: FT Numeris
DUTCH	Netherlands ITR6 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

Loopback support for IDSL card (BRI/LT)

The Ascend ISDN Digital Subscriber Line (IDSL) card now supports loopback tests to diagnose problems in the IDSL connection.

Overview

Previously, the IDSL card (also known as the BRI/LT slot card) for the MAX only supported one loopback command; namely, the a loopback at the remote TA or Pipeline. Now it can command any device on the IDSL connection to loop back towards the MAX. Specifically, you can command the remote TA or Pipeline to loop back the signal from the IDSL card back toward the IDSL card, or loop back at any intermediate repeater (see Figure 7).




Figure 7. IDSL connection with repeaters

In Figure 7, you could set up a loopback test from the MAX to any of the ISDN repeaters, or from the MAX all the way to the remote ISDN at the end of the connection. This allows you to isolate trouble over the entire connection.

Configuring a loopback test

To configure a loopback test on the BRI lines provided by the IDSL slot card:

- 1 Select BRI/LT > Line Diag > Line *N*, where *N* is the number of the line you want to loopback.
- 2 Specify the EOC Address of the device that is the loopback point for the test.
 - 0 specifies the remote TA or Pipeline provides a loopback returning all signals from the MAX to the MAX
 - 1 specifies a loopback at the repeater nearest the MAX
 - 2-6 specifies a loopback at the next nearest, etc. repeater
 - 7 specifies all devices loop back toward the MAX
- 3 Select Line Loopback and press Enter.
- 4 In the confirmation dialog that appears, select 1=Line *N* LB.
While the line is being looped back, normal data transfer is disrupted.
- 5 Press Escape to cancel the loopback.

New parameters

This section described the new items have been added to the BRI/LT Line Diag menu.

EOC Address

Description: Specifies the Embedded Operations Channel (EOC) address from of the device IDSL line that the MAX commands to provide a loopback.

Usage: Specify one of the following values:

- 0 (the default) addresses the loopback command to the remote TA or Pipeline. All signals from the MAX are rolled back at the remote TA or Pipeline to the MAX.
- 1 addresses the loopback command to the repeater nearest the MAX. All signals from the MAX are rolled back at the repeater to the MAX.
- 2-6 addresses the loopback command to the next nearest, etc. repeater.
- 7 addresses the loopback command to all devices on the IDSL line.

Note: The EOC address setting reverts to its default value of 0 whenever you exit the Line Diag submenu.

Location: BRI/LT > Line Diag > line *n*

Sealing Current

Description: Sealing Current allows you to enable “sealing” on the loop. Sealing refers to the ability of the IDSL card to send some current (40V) on the line when enabled. You typically use this feature to keep the physical connection from corroding. This could occur if there is no activity on the line such as when there is no device connected on the other end.

Usage: Specify Yes to enable sealing. The default value is Off.

Dependencies: Note that the Sealing Current setting is not saved to the MAX permanent memory. This means that whenever you reboot the MAX, the Sealing Current parameter reverts to its default value of 0.

Location: BRI/LT > Line Diag > line *n*

User interface changes

The individual commands in the BRI/LT Line Diag menu are no longer numbered. In previous releases, the menu appeared as follows:

```
40-200 Line Diag
40-20X Line X...
40-201 Line LoopBack
40-202 Corrupt CRC
40-203 UnCorrupt CRC
40-204 Rq Corrupt CRC
40-205 UnRq Corrupt CRC
40-206 Clr NEBE
40-207 Clr FEBE
```

The new menu is presented below:

```
40-200 Line Diag
40-20X Line X...
>EOC Address=
Line LoopBack
Corrupt CRC
UnCorrupt CRC
Rq Corrupt CRC
UnRq Corrupt CRC
Clr NEBE
```

```
Clr FEBE
Sealing Current
```

New status messages

When you enable the Loop Sealing Current, the following message appears in the edit window:

```
Message #242
  Loop Sealing Current
    now ON
```

When you disable the Loop Sealing Current, the following message appears in the edit window:

```
Message #243
  Loop Sealing Current
    now OFF
```

MAX supports new CSLIP Auto Detect parameter

Previously, when you brought up a SLIP session, compression was the default. Now, you can bring up a SLIP session and choose no compression until you receive a VJ Compressed CSLIP packet. At that point, the MAX switches automatically to VJ compression mode.

CSLIP Auto Detect

Description: Enables and disables auto-detect of VJ Compressed CSLIP packet. Previously, when you brought up a SLIP session, compression was the default. Now, you can bring up a SLIP session and choose no compression until you receive a VJ Compressed CSLIP packet. At that point, the MAX switches automatically to VJ compression mode.

Usage: The CSLIP Auto Detect parameter has two options:

- Yes: VJ Compression is always on for all CSLIP packets.
- No: Compression is off for CSLIP packets until the MAX receives a VJ Compression CSLIP packet. When this occurs, the MAX starts VJ compression of all subsequent CSLIP Packets. This is the default.

Note: Depends on VJ Comp parameter. Applies only if VJ Comp=Yes, otherwise CSLIP Auto Detect=NA.

Example: CSLIP Auto Detect=Yes

Location: Mod Config > Tserv options

More information provided for SLIP connections

You can now specify the kind of information the MAX reports when a user connects over a SLIP link.

You can now specify the information that the MAX reports to users when they establish a Serial Line Internet Protocol (SLIP) connection.

Previously, the MAX always reported the following information whenever a user connected:

```
Entering SLIP Mode
IP address is 192.1.1.1
MTU is 1500
```

Below is an example of the kind of information the MAX can now report:

```
Entering SLIP Mode
IP address is 192.1.1.1
MTU is 1500
Netmask: 255.255.255.0
Gateway: 192.168.6.181
```

The Netmask label identifies the subnet mask the MAX is using. The Gateway label identifies the MAX unit's IP address. The sections that follow describe these parameters.

New parameters

The MAX interface includes three new parameters to support this feature:

- SLIP Info
- IP Netmask Msg
- IP Gateway Adrs Msg

These parameters are described below.

SLIP Info

Description: Specifies the type of information the MAX reports to SLIP users.

Usage: Specify one of the following values:

- Basic (the default)
Specifies that the MAX only reports the SLIP user's IP address and the Maximum Transmission Unit (MTU).
- Advanced
Specifies that the MAX reports the SLIP user's IP address, the MTU, the Netmask, and the Gateway to SLIP users. Note that the gateway is the MAX unit's IP address.

Location: Ethernet>Mod Config>TServ Options

IP Netmask Msg

Description: Specifies the text the MAX displays before the netmask field in the SLIP session startup message.

Usage: Specify a text message. You can enter up to 64 characters. The default is Netmask:.

Dependencies: Keep this additional information in mind.

- IP Netmask Msg does not apply unless you set SLIP Info to Advanced.

Location: Ethernet>Mod Config>TServ Options

IP Gateway Addr Msg

Description: Specifies the text the MAX displays before the MAX IP address field in the SLIP session startup message.

Usage: Specify a text message. You can enter up to 64 characters. The default is Gateway:.

Dependencies: Keep this additional information in mind.

- IP Gateway Addr Msg does not apply unless you set SLIP Info to Advanced.

Location: Ethernet>Mod Config>TServ Options

Multilink or MP+ call now can span multiple MAX units

Multiple MAX units can now be configured to form a stack, or group of MAX units, that allows a Multilink PPP (MP) or MP+ call to span the MAX units in the stack. This feature, previously added to the MAX 4000, is now available for other MAX units.

Call spanning using a stack configuration can be effective when:

- A MAX running MP+ is asked for another phone number, and has no available lines
- A rotary hunt-group uses the same phone number to access multiple MAX units, making it impossible to assume that a subsequent call is answered by the same MAX as the original call

MP/MP+ call spanning is protocol independent, and should work with all protocols supported by the MAX.

Note: Stacking requires any MP caller to use the MP endpoint discriminator. The same is true of MP+. All Ascend products and most other products that support MP or MP+ use an endpoint discriminator, but the specification for MP does not require it.

How MP/MP+ call spanning works

A stack is a group of MAX units that have the same stack information, and are on the same physical LAN. There is no *master* MAX; the MAX units in the stack use an Ethernet multicast packet to locate each other.

WAN access features

Multilink or MP+ call now can span multiple MAX units

Multicast packets usually cannot cross a router, so the MAX units in a single stack must be on the same physical LAN. MAX units running in a stack can generate fairly high levels of network traffic, which is another reason to keep them on the same physical LAN.

Bundle ownership

Although MAX stacks do not have a master MAX, each MP/MP+ bundle has a bundle owner. The MAX that answers the first call in the MP/MP+ bundle is the *bundle owner*. If a bundle spans more than one MAX in a stack, an exchange of information flows between the MAX units in the bundle.

Stacking requires an endpoint discriminator. Every MP/MP+ call that comes to any member of the stack is compared to all existing MP/MP+ calls in the MAX stack to determine whether it is a member of an existing bundle. If the call belongs to an existing bundle, the MAX that answered and the bundle owner exchange information about the bundle. Furthermore, the MAX that answered the call forwards all incoming data packets over the Ethernet to the bundle owner.

Outgoing data

To balance the load among all available WAN channels, outgoing data packets for the WAN are assigned to available channels in a bundle on a rotating basis. If an outgoing packet is assigned to a channel that is not local to the bundle owner, the bundle owner forwards the packet over the Ethernet to the MAX that owns the non-local channel.

Real and stacked channels

For the purpose of this description, *real* channels are those channels that connect directly to the MAX that owns the bundle. *Stacked* channels connect to a MAX that transfers the data to or from the MAX that owns the bundle.

For example, assume the initial call of an MP/MP+ bundle connects to MAX #1. This connection is a *real* channel. Next, the second call of the bundle connects to MAX #2. This connection is a *stacked* channel. MAX #1 is the bundle owner, and it manages the traffic for both channels of the bundle. MAX #2 forwards any traffic from the WAN to MAX #1, for distribution to the destination. See Figure 8.




Figure 8. Packet flow from the slave channel to the Ethernet

Note: This graphic does not illustrate traffic from the master MAX. WAN traffic received on the master channel by MAX#1 is forwarded directly to the destination.

Likewise, MAX#1 receives all Ethernet traffic destined for the bundle, and disperses the packets between itself and MAX#2. See Figure 9. MAX#1 forwards some of the packets across the WAN through a real channel. MAX#2 sends the rest of them through a stacked channel.




Figure 9. Packet flow from the Ethernet

Connection profiles not shared within a stack

A stack does not support sharing of local Connection profiles between the MAX units in the stack. Every MAX in the stack that is set up to use internal authentication must retain all authentication information for every call. You can eliminate this requirement by using a centralized authentication server, such as RADIUS.

Phone numbers for new MP+ and MP-with-BACP channels

When a MAX has to add a channel for a MP+ or MP-with-BACP call, it provides a local phone number for the new channel. However, sometimes the MAX that answers the call cannot provide a local phone number for the additional channel, because all the channels that connect directly to it are busy. In that case, the MAX requests other members of the stack to supply a phone number for the additional channel.

An MP call does not pass phone numbers when it adds a channel. The originator of the call must know all of the possible phone numbers to begin with.

If each MAX in the stack is accessed through a different phone number, the originator of the call must know all of the possible phone numbers. An alternative in this instance is to use BACP or MP+ to obtain the phone number of a MAX with a free channel.

Performance considerations for MAX stacking

There is no limit to the number of *stacked* channels in single call or in a stack of MAX units, other than the limit for each individual MAX. The MAX 4000, MAX 2000, and MAX 1800 each support up to 40 stacked channels. The MAX 200 Plus supports up to three stacked channels. A MAX can handle n real channels and $n/3$ *stacked* channels.

There is no theoretical limit to the number of MAX units in a stack, other than performance considerations. Since all data from stacked channels crosses the LAN, performance could suffer with a large number of MAX units in the stack and many stacked channels in use.

WAN access features

Multilink or MP+ call now can span multiple MAX units

Performance overhead increases when stacked bundles span multiple boxes. In a bundle of 6 channels, 4 of which are real and 2 are stacked, the overhead is the actual bandwidth of the two stacked channels ($2 \times 64 = 128\text{K}$). The actual payload data of the 6 channels with a 2:1 data compression is $6 \times 2 \times 64 = 768\text{K}$. The overhead is 128 over 768, or 16%. In a two-channel bundle with one real and one stacked channel, with the same compression, the overhead is 25%.

Take into account that you do not know ahead of time how many bundles will span the stack, or how many multi- or single-channel calls you are going to get. You can base an estimate on your traffic expectations. But in most situations, the majority of bundles will be on a single MAX, for which there is no overhead.

Suggested LAN configurations

Calculations like the ones mentioned above show that when your MAX stack handles 82 single-channel calls, 41 two-channel stacked calls, and 41 two-channel nonstacked calls, the Total Ethernet usage is approximately 5116Kbps. Since Ethernet capacity generally does not achieve more than 50% utilization, this configuration uses up the available Ethernet bandwidth.

The total number of channels in this configuration is 246. Therefore, a stack of three MAX units, each having three T1 lines with this usage profile, utilizes all of the Ethernet bandwidth.

The basic limitation from the above examples is the speed of the LAN. One way to increase the speed of your LAN is to attach each MAX to a separate port of a 10/100 Ethernet switch, then use a 100Mbps connection to the backbone LAN. This allows each MAX to utilize up to a full 10Mb Ethernet and the entire stack combined can generate up to full 100Mb of Ethernet data. Once again assuming that the 100Mbps is saturated at 50% usage, we can now use up to 51200Kbps of bandwidth, or 10 times more than in the example above. Note that the success of this strategy depends on limiting stacked channels per MAX to the $n/3$ limit mentioned above.

Suggested hunt-group configurations

Whenever you have MAX units in a stack, it is important to limit the number of multichannel calls that are split between the MAX units. The following suggested configurations reduce the overhead for a multichannel call by keeping as many channels as possible on the same MAX.

MP+ and MP-with-BACP calls

Figure 10 shows the suggested hunt-group setup for a typical MAX stack that receives only PPP, MP+, or MP-with-BACP calls. Each MAX has three T1 lines. All the T1 lines in a MAX share a common phone number and they are in a hunt-group that does not span MAX units. The illustration shows these three local hunt-groups with phone numbers 555-1212, 555-1213, 555-1214. In addition, a global hunt-group, 555-1215 spans all the T1s of all the MAX units in the stack.

Users that access the MAX dial 555-1215, the global hunt-group number. The telephone company has set up the global hunt-group to distribute incoming calls equally among the MAX units. Namely, the first call dialing 555-1215 goes to MAX#1, the second call to MAX #2, and so on. If you use this configuration, you must configure each of the MAX unit's Line profiles with the local hunt-group numbers. For example, for MAX #1 in Figure 10, you would set the Ch n # parameters to 12 (the last two digits of the 555-1212 hunt-group number).

You can achieve the same distribution without a global hunt-group by having one third of the users dial 555-1212, one third dial 555-1213, and one third dial 555-1214. You can leave the Ch *n* # parameters at their default setting (null) if you do not have a global hunt-group.




Figure 10. Hunt-groups for a MAX stack handling both MP and MP+ calls

Viewing Figure 10, suppose an MP+ call is connected to MAX #1. When that call needs to add a channel, it requests an add-on number from the MAX, and the MAX returns 12 (for 555-1212) as long as a channel in the local T1 lines is available. This means the bundle will not span multiple MAX units as long as a channel is available in the local hunt-group.

The Figure 10 configuration tends to break down if MAX units receive MP-without-BACP calls. Spreading the calls across the MAX stack (by dialing the global hunt-group) results in the worst possible performance, because MP-without-BACP must know all of the phone numbers before the caller places the first call.

MP-without-BACP calls

Figure 11 shows a site that supports only MP-without-BACP calls. For this site, the telephone company has set up a global hunt-group that first completely fills MAX #1, then continues to MAX #2, and so on. This arrangement tends to keep the channels of a call from being split across multiple MAX units, keeping overhead low.




Figure 11. Hunt-groups for a MAX stack handling only MP-without-BACP calls

MP+ calls and MP calls with or without BACP

For a MAX that receives MP+ calls and MP calls with or without BACP, you can use a configuration similar to the one shown in Figure 10. In this case, however, you set up the global hunt-group differently than explained in “MP+ and MP-with-BACP calls.” You set up the global hunt-group to help prevent MP-without-BACP calls from being split across multiple MAX units in the stack. As in “MP-without-BACP calls,” calls dialing 555-1215 first completely fill the channels of MAX #1, then continues to MAX #2, and so on.

Both MP+ and MP callers dial the global hunt-group number to connect to the stack. Channels added to the MP+ and MP bundles are handled as explained in “MP-without-BACP calls,” and “MP+ calls and MP calls with or without BACP.” Be sure to set the Ch n # parameters as explained in “MP+ calls and MP calls with or without BACP.”

MP+ and MP-with-BACP callers do not have to dial the global hunt-group numbers to connect. Only the MP-without-BACP callers need to dial the global hunt-group. You can achieve an even distribution of MP+ and MP-with-BACP calls by having one third dial 555-1212, one third dial 555-1213, and one third dial 555-1214. You can leave the Ch n # parameters at their default setting (null) in this situation.

Configuring a MAX stack

To configure a MAX stack, proceed as follows for each MAX in the stack:

- 1 Open the Ethernet > Mod Config menu, and select Stack Options, as shown in the following sample menu:

```
90-A** Mod Config
RADIUS Server
Log
ATMP
Modem Ringback=Yes
AppleTalk
SNTP Server
>Stack Options...
UDP Checksum=No
```

When you press Enter, the Ethernet > Mod Config > Stack Options menu appears. For example:

```
90-A** Mod Config
>Stack Options...
Stack Enabled=Yes
Stack Name=astack
UDP Port=5151
```

- 2 Set Stack Enabled to Yes (Stack Enabled=Yes).
- 3 Set the Stack Name parameter to a unique name for the stack.

A stack name is 16 characters or less. This is the name members of a stack use to identify other members of the same stack. The stack name must be unique among all MAX units that communicate with each other, even if they are not on the same LAN.

If a MAX receives calls from two MAX units on different LANs, and the two units are members of different stacks with the same stack name, the MAX receiving the calls assumes the two MAX units with the same stack name are in the same bundle.

Note: Multiple stacks can exist on the same physical Ethernet LAN if the stacks have different names.

4 Specify the UDP port.

This is a reserved UDP port for intrastack communications. The UDP port must be identical for all members of a stack, but is not required to be unique among all stacks.

Disabling a MAX stack

To disable a stack, specify Stack Enabled=No for each of the MAX units in the stack.

Adding and removing a MAX

You can add a MAX to an existing stack at any time without rebooting the MAX or affecting stack operation. Since a stack is a collection of peers, none keeps a list of the stack membership. The MAX units in a stack communicate when they need a service from the stack.

Removing a MAX from a stack requires care, because any calls using a channel between the MAX to be removed and another MAX in the stack could be dropped. There is no need to reboot a MAX removed from a stack.

Parameter Reference

This release adds two new parameters: Stack Enabled and Stack Name.

Stack Enabled

Description: Stack Enabled enables MP and MP+ call spanning for the MAX. When stack Enabled=Yes, a *stack*, or group of MAX units that have the same stack information and are on the same physical LAN. The MAX units in the stack use an Ethernet multicast packet to locate each other. Once a stack is created, every MP/MP+ call that comes to any member of the stack is compared with MP/MP+ calls to other members of the stack to determine if it is part of an already existing bundle.

If you disable this parameter, all channels in an MP or MP+ bundle must exist on a single MAX. If you enable this parameter, MP and MP+ bundles can span any of the MAX units in the stack. Note that MP and MP+ bundles that span MAX units cause additional traffic on the Ethernet as the MAXs in the bundle route packets between them.

Stacking requires an endpoint discriminator. Every MP/MP+ call that comes to any member of the stack is compared to all existing MP/MP+ calls in the MAX stack to determine whether it is a member of an existing bundle. If the call belongs to an existing bundle, the MAX that answered and the bundle owner exchange information about the bundle. Furthermore, the MAX that answered the call forwards all incoming data packets over the Ethernet to the bundle owner.

Usage: Select Yes or No.

- Yes enables MP and MP+ call spanning.
- No disables MP and MP+ call spanning.

Location: Ethernet > Mod Config > Stack Options

See Also: Stack Name

Stack Name

Description: Stack Name defines a unique name for a MAX stack. This is the name members of a stack use to identify other members of the same stack. The stack name must be unique among all MAX units that communicate with each other, even if they are not on the same LAN.

If a MAX receives calls from two MAX units on different LANs, and the two units are members of different stacks with the same stack name, the MAX receiving the calls assumes the two MAX units with the same stack name are in the same bundle.

Multiple stacks can exist on the same physical Ethernet LAN if the stacks have different names.

Usage: Press Enter to open a text field, then type a unique name for the MAX stack, up to a maximum of 16 characters.

Location: Ethernet > Mod Config > Stack Options

See Also: Stack Enabled

Updated Microsoft Callback Control Protocol support

This release adds support for Microsoft's CallBack Control Protocol (CBCP). CBCP is a Link Control Protocol (LCP) option negotiated at the beginning of Point to Point Protocol (PPP) sessions. CBCP authenticates callers by means of user names and passwords, and offers additional security to enable the MAX to ensure that connections are to known users.

Introduction

Microsoft developed CBCP to address a need for greater security with PPP connections. The standardized callback option defined in RFC 1570 has a potential security risk because the authentication is performed after the callback. CBCP callback like Ascend's proprietary callback, occurs after authentication, leaving no potential security hole.

CBCP also offers features not available with the standard callback defined in RFC 1570. The client side supports a configurable time delay to allow users to initialize modems or enable supportive software before the MAX calls the client. You can configure the MAX with a phone number to use for the callback, or you can configure it to allow the client to specify the phone number used for the callback.

Currently, Microsoft's Windows NT 4.0 and Windows 95 software support client-side authentication using CBCP. The MAX now supports a CBCP central-site solution.

Ascend's implementation of CBCP

CBCP is an option negotiated during the LCP negotiation of a PPP session. While support for CBCP is configured systemwide on the MAX, not every connection must negotiate its use. Parameters have been added to the Answer Profile under Ethernet > Answer > PPP Options, and to each Connection Profile under Ethernet > Connections > Encaps Options. The calling and called sides of a PPP session initiate authentication after acknowledging that CBCP is to be used.

Note: Currently, the MAX does not initiate LCP negotiation of CBCP. The MAX responds to *caller* requests to configure CBCP.

The MAX employs the user name and password to link a caller with a specific Connection profile or RADIUS User profile. Configured CBCP parameters in that Connection profile specify variables for the callback. If, at any point, the client and the MAX disagree about any CBCP variables, the MAX might drop the connection.

Both sides of the connection must agree on whether the callback phone number is supplied by the client or by the MAX. A new trunk group parameter, configured on the MAX, supplies a trunk group that is prepended to phone numbers when supplied by the client.

Negotiation of CBCP

Following are the steps from initial connection to MAX callback:

- 1 Caller connects to MAX.
- 2 LCP negotiations begin.
Caller and MAX must agree to use CBCP. Otherwise, the MAX terminates the connection.
- 3 After successful LCP negotiation, both sides have acknowledged that CBCP will be used, and CBCP begins after authentication.
- 4 Caller authenticates itself to MAX. If authentication fails, the MAX terminates the connection.
- 5 The MAX verifies that the profile has CBCP Mode set. CBCP begins.
- 6 The MAX sends a request to determine if a callback is to occur. The caller's configuration must match the CBCP Mode value on the MAX.
The client also supplies to the MAX the number of seconds it should delay before initiating the callback, and, if applicable, the phone number.
- 7 If both sides agree on which phone number the MAX will dial, the client clears the connection.
- 8 The MAX delays the callback on the basis of the previous negotiation.
- 9 The MAX dials the client, by applying information from the same profile used in previous negotiation.

Configuring Microsoft's CBCP to use a Connection Profile

To configure CBCP to work with a Connection profile:

- 1 Open the Ethernet > Answer > PPP Options menu.
- 2 Set CBCP Enable = Yes.

- 3 Open the Ethernet > Connections > *Any Connection profile* > Encaps Options menu.
- 4 Set CBCP Mode to the callback mode to be offered the caller.
- 5 If the caller is supplying the phone number, set CBCP Trunk Group to the value (4-9) that the MAX prepends to the number when calling back.
- 6 Save your changes.

New parameters

The following parameters have been added to the VT100 interface:

CBCP Enable

Description: Specifies how the MAX responds to caller requests to support CBCP.

Usage: Press Enter to cycle through the choices.

- Yes specifies the MAX will positively acknowledge, during LCP negotiations, support for CBCP.
- No specifies the MAX will reject any request to support CBCP.
No is the default.

Location: Ethernet > Answer > PPP Options

See Also: CBCP Mode, CBCP Trunk Group

CBCP Mode

Description: Specifies what method of callback the MAX offers the incoming caller.

Usage: Press Enter to cycle through the choices. You can specify one of the following settings:

Setting	Description
No Cback	Applies for Windows NT or Windows 95 clients who must not be called back. Because CBCP has been negotiated initially, the Windows clients must have validation from the MAX that no callback is used for this connection.
User Num	Specifies that the caller will supply the number the MAX uses for the callback.
Prof Num	Specifies the MAX will use the number in Ethernet > Connections > <i>Any Connection profile</i> > Dial # for the callback
User Num or No Cback	Specifies that the caller has the option of either supplying the number to dial or specifying that no callback is used for the call. If no callback is chosen, the call will not be disconnected by the MAX.

Dependencies: CBCP Mode applies only if CBCP is successfully negotiated for a connection. Encaps=PPP or MPP or MP.

Location: Ethernet > Connections > *Any Connection Profile* > Encaps Options

See Also: CBCP Enable, CBCP Trunk Group

CBCP Trunk Group

Description: Assigns the callback to a MAX trunk group. This parameter is used only when the caller is specifying the phone number the MAX uses for the callback. The value in CBCP Trunk Group is prepended to the caller-supplied number when the MAX calls back.

Usage: Press Enter to open a text field. Then type a number from 4 to 9. The default is 9.

Dependencies: CBCP Trunk Group applies only if CBCP is negotiated for a connection. Encaps=PPP or MPP or MP.

Location: Ethernet > Connections > *Any* Connection Profile > Encaps Options

See Also: CBCP Enable, CBCP Mode

Configuring Microsoft's CBCP to use a RADIUS Profile

New RADIUS attributes support CBCP in User profiles. Ascend-CBCP-Enable specifies how the MAX responds to caller requests to support CBCP. Ascend-CBCP-Mode specifies the method of callback the MAX offers the incoming caller. Ascend-CBCP-Trunk-Group assigns the callback to a MAX trunk group.

Note: Make sure you set CBCP Enable=Yes in the Ethernet > Answer > PPP Options menu.

Ascend-CBCP-Enable (112)

Description: Specifies how the MAX responds to requests by callers to support CBCP.

Usage: Specify one of the following settings:

- CBCP-Enabled (0)—Specifies that the MAX will positively acknowledge, during LCP negotiations, support for CBCP.
- CBCP-Not-Enabled (1)—Specifies that the MAX will reject any request to support CBCP.

See Also: Ascend-CBCP-Mode, Ascend-CBCP-Trunk-Group

Ascend-CBCP-Mode (113)

Description: Specifies what method of callback the MAX offers the incoming caller.

Usage: Specify one of the following values:

- CBCP-No-Callback (1)—Applies for Windows NT or Windows 95 clients who must not be called back. Because CBCP has been negotiated initially, the Windows clients must have validation from the MAX that no callback is used for this connection.
- CBCP-User-Callback (2)—Specifies that the caller will supply the number the MAX uses for the callback.
- CBCP-Profile-Callback (3)—Specifies that the MAX will use the number in Ascend-Dial-Number for the callback

- CBCP-User-Or-No (7)—Specifies that the caller has the option of either supplying the number to dial or specifying that no callback is used for the call. If no callback is chosen, the call will not be disconnected by the MAX.

Dependencies: Ascend-CBCP-Mode applies only if CBCP is successfully negotiated for a connection.

See Also: Ascend-CBCP-Enable, Ascend-CBCP-Trunk-Group

Ascend-CBCP-Trunk-Group (115)

Description: Assigns the callback to a MAX trunk group. This attribute is used only when the caller is specifying the phone number the MAX uses for the callback. The value in Ascend-CBCP-Trunk-Group is prepended to the caller-supplied number when the MAX calls back.

Usage: You can specify a number between 4 and 9, inclusive. The default is 9.

Dependencies: Ascend-CBCP-Trunk-Group applies only if CBCP is negotiated for a connection.

See Also: Ascend-CBCP-Enable, Ascend-CBCP-Mode

MAXDial/IP support for immediate modem call restriction

MAXDial now supports the call restriction modes for immediate modem service. The immediate modem service in the MAX terminal server now supports call restriction modes. MAXDial has been updated to connect to a MAX that has the updated immediate modem service.

In the MAX system software, immediate modem service now supports three authentication modes: "none", "global password", or "user". When the immediate modem service is configured in "user" mode, the MAXDial software must be configured to supply a user name as well as a password to use the digital modems for dialing out. To implement this change, the MAXDial Ports Control Panel has a new User Name field.

If the immediate modem service is configured in the "none" or "global password" mode, the User Name field may be empty. If the immediate modem service is configured in "user" mode, MAXDial displays a "login name required" error message if its User Name field is empty and the user attempts to dial out on a digital modem.

Configure Port

MaxDial Serial Port (COM5)

☐ Unassigned

☒ Assign this port to an Ascend MAX

Name: techpubs-lab-20

IP Address: 192.168.1.101 Find...

Immediate Modem Port: 5000

User name:

Password:

Test Connection...

OK Cancel

Support added for multiple host selection

You can now specify up to three authentication hosts for Defender authentication, so that these hosts can serve as backups for each other. Previously, you could only specify one authentication host, although the Defender authentication mechanism itself allowed multiple hosts.

How it works:

There are three major stages in authentication using AssureNet Pathways' Defender. The MAX' behavior will depend upon the stage the call dialing the MAX was in when the connection with the host is lost.

WAN access features

Support added for multiple host selection

Table 12. Token card authentication

Stage	Description	MAX Behavior at this stage
1	<p>Usually a short time after the caller has connected to the MAX and before the MAX has received the first prompt from the authentication host.</p> <p>The Defender server provides the text of the prompts or challenges, and the MAX passes them through to the caller.</p>	<p>Calls in Stage 1 are preserved if an authentication host is unavailable or loses its connection.</p> <p>This might be the case when the very first caller is authenticating with Defender after the router boots up, and the first authentication host is unavailable. The Defender authentication code in the router will try the second and third hosts in order to authenticate the user.</p>
2	<p>During the time the caller is interacting with the authentication host, but before the authentication sequence is complete.</p> <p>The Defender uses a challenge-response protocol, with a token card to provide the responses.</p>	<p>Calls in Stage 2 are never preserved if an authentication hosts loses its connection.</p> <p>Defender has no mechanism for having one authentication server take over for another if the first loses connection in the middle of a state.</p>
3	<p>When the caller has completed authentication and is interacting with the MAX normally (either asynchronously or framed).</p>	<p>Callers in Stage 3 are not dropped by the router since their calls are already authenticate. However, because the host on which they authenticated is no longer available, their logout time will not be sent (as would be the case if the host had remained connected).</p> <p>Defender provides no mechanism to notify one authentication host when a user call that was authenticated by another host is terminated.</p>

When no authentication host is available

When a MAX can not establish contact with any of the authentication hosts in the list, all sessions are dropped, including calls in Stage 1.

If a caller who has been disconnected tries again to make a connection, the MAX will begin again the process of connecting to authentication hosts on the list until it either succeeds or has tried every host in the list.

User interface changes

You can specify up to three authentication hosts in the Auth submenu of the Ethernet Configuration Profile using the Auth Host #1, #2, and #3 parameters. Previously these were N/A when Auth-DEFENDER. You can one of the Auth Host parameters, as shown in the example for Auth Host #2. The authentication process checks for null IP address and does not attempt to use a null IP address.

```
90-100 Mod Config
Auth
>Auth=Defender
Auth Host #1=137.175.80.62
Auth Host #2=0.0.0.0
Auth Host #3=137.175.80.24
Auth Port=2626
Auth Src Port=0
Auth Timeout=30
Auth Key=
```

Changes to syslog messages

This feature introduces a set of syslog messages reporting the status of the Defender authentication subsystem. The new syslog messages are reported with "LOG_DEFAULT" and "LEVEL_INFO" priorities. The following lines exemplify the new syslog messages.

```
Nov 14 15:59:34 137.175.85.20 ASCEND: AuthHost 137.175.81.24
Activated
Nov 14 15:51:10 137.175.85.20 ASCEND: AuthHost 137.175.81.24
Fails auth
Nov 14 15:51:10 137.175.85.20 ASCEND: AuthHost 137.175.80.24
Refuses connect
Nov 14 16:03:05 137.175.85.20 ASCEND: AuthHost 137.175.81.24
Closed connection
Nov 14 16:05:59 137.175.85.20 ASCEND: AuthHost 137.175.81.24
Address Changed
Nov 14 16:06:31 137.175.85.20 ASCEND: AuthHost 137.175.81.24
New Authmethod
```

Message format

All Defender syslog messages report the standard Ascend header plus Defender-formatted detail: "AuthHost xx.yy.zz.aa Statusx" where "Statusx" has the values shown in the following table:

Table 13. Status indicators

Status	Description
Activated	A Defender Host has been found and a connection successfully established. This state is reported when an authentication session is active and ready to authenticate.
Fails auth	A Defender Host has been found, but the router and the Defender authentication server do not agree on their mutual authentication key.
Refuses connect	The host either is not responding at all or has no active Defender Server running.
Closed connection	An active Defender authentication server has ended its connection. This would reflect either a failure of the server software or explicit request by an administrator for the server to shutdown.
Address Changed	The MAX administrator has changed the IP address, the port number, or the authentication key of the active authentication server. This forces the Defender authentication subsystem to close its connection with the active server and start searching for a new one.
New Authmethod	The MAX administrator has changed the authentication method from "DEFENDER" to something else, causing the Defender authentication subsystem to break an active connection.

Limiting terminal server access per user

This feature allows the MAX to limit particular users to a subset of terminal server commands.

The Framed Only parameter in the Answer profile and the Connection profiles allows administrators to limit particular users to the PPP, SLIP, CSLIP, and Quit commands in the MAX terminal server interface.

Configuring per-user access to terminal server commands

You can configure per-user access to the terminal server commands in the Answer profile or in the Connection profile:

- The Answer profile affects users who do not have a Connection profile, users with a Name/Password profile, or RADIUS-authenticated users whose connections are built in part with the Answer profile
- The Connection profile only affects individual users connecting to the MAX using a particular Connection profile

To configure per-user access to the terminal server:

- 1 Select Ethernet > Answer > Session Options *or*
Ethernet > Connections > a *Connection profile* > Session Options
- 2 Specify one of the following values for Framed Only:
 - No (the default)
Specifies that terminal server users connecting through this profile have unlimited access to the terminal server commands.
 - Yes
Specifies that terminal server users connecting through this profile only have access to the PPP, SLIP, CSLIP, and Quit terminal server commands.
- 3 Save and exit the profile.

If a user restricted to these commands tries to execute any other terminal server command, the MAX displays the following message:

Unauthorized Terminal Server Command.

Parameter reference

This section describes the new MAX parameter.

Framed Only

Description: Specifies whether the user is allowed access to all the terminal server commands or to a subset of them.

Usage: Specify one of the following values:

- No (the default)
Specifies that terminal server users connecting through this profile have unlimited access to the terminal server commands.
- Yes
Specifies that terminal server users connecting through this profile only have access to the PPP, SLIP, CSLIP, and Quit terminal server commands.

Dependencies: Keep this additional information in mind:

- Framed Only has no affect if TS Enabled is set to No in the Ethernet > Mod Config > TServ Options submenu.

WAN access features

User-configurable call blocking after failed connection attempt

- PPP, SLIP, and CSLIP must be enabled in the Ethernet > Mod Config > TServ Options submenu before users can start a PPP, SLIP, or CSLIP session.

Location: Ethernet > Answer > Session Options

Ethernet > Connections > *any Connection profile* > Session Options

User-configurable call blocking after failed connection attempt

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. Previously you could not automatically stop the Ascend unit from attempting to place an outgoing call on a connection that repeatedly fails.

Overview

When an Ascend unit attempts to make a connection and the attempt fails, the Ascend unit continues to attempt to complete the connection. The number of retry attempts allowed without using this new call blocking feature is very large; successive retries can cause excessive charges, congestion, and performance problems. This feature enables you to specify the number of unsuccessful attempts to place a call that an Ascend unit 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 Ascend unit continues to block further calls for a the period of time you specify.

Configuring call blocking

- 1 Open the Session options submenu of the Connection Profile.
- 2 Select Block calls after and enter the number of retry attempts to allow the Ascend unit to make when placing a call.
- 3 Select Blocked duration and specify the length of time during which the Ascend unit will continue to block calls to number in the Connection Profile.

Note: This feature applies only to outgoing calls that are not answered by the far end. It does not apply to:

- incoming calls
- outgoing calls that connect and are immediately disconnected

Parameter reference

Two new parameters have been added to the Session submenu of the Connection Profile.

Block calls after

Description: Specifies how many unsuccessful attempts the Ascend unit will make before beginning to block outgoing calls.

Usage: Enter the number of connection attempts permitted before the Ascend unit blocks calls for the connection. The maximum number you can enter is 65535 (65535 attempts). The default is 0.

Location: Session Options submenu of the Connection Profile.

See Also: Blocked duration

Blocked duration

Description: Specifies the length of time in seconds during which the Ascend unit will block outgoing calls.

Usage: Enter the number of seconds for the Ascend unit to block all calls made to the connection. When this period has elapsed, the unit will again allow calls to this connection.

Location: Session Options submenu of the Connection Profile.

See Also: Block calls after

RADIUS attributes added

Two new attributes corresponding to the new parameters have been added to the RADIUS dictionary to enable this feature using RADIUS:

- Ascend-Call-Attempt-Limit (123)
This attribute has the same function as the Block Calls After parameter, described above.
- Ascend-Call-Block-Duration (124)
This attribute has the same function as the Blocked Duration parameter, described above.

Specifying Shared Profiles per Connection Profile

Previously, you could configure shared profiles on a per-MAX basis by setting Ethernet > Mod Config > Shared Prof = Yes. This release adds a Shared Prof parameter to the Connection profile, enabling you to allow users to share specific profiles. You can also set a new RADIUS attribute to allow multiple incoming users to share a user file.

To use the new parameter or the new attribute, you must disable profile sharing on a per-MAX basis. That is, if you specify:

Ethernet > Connections > *Any Connection Profile* = Yes
or, in a RADIUS user profile:

Ascend-Shared-Profile-Enable = Shared-Profile-Yes

You must also specify:

Ethernet > Mod Config > Shared Prof = No

Shared Prof

Description: Enables multiple incoming callers to share a local Connection profile. Note that to apply shared profiles on a per-Connection-profile basis, you have to disable profile sharing on a system-wide basis by setting Ethernet > Mod Config > Shared Prof = No.

Usage: Press Enter to toggle between Yes and No.

- Yes specifies that multiple incoming calls can share a local Connection profile.
The MAX must first authenticate the caller by applying the profile's Name and Recv PW parameters. If an incoming call has an IP address that conflicts with an existing caller IP address, or if the MAX would have to assign a conflicting IP address from the IP address pool, the MAX rejects the call.
- No specifies that multiple incoming calls cannot share a local Connection Profile.
No is the default.

Dependencies: Shared Prof for Connection profiles applies only if you have disabled shared profiles for the MAX as a whole with Ethernet > Mod Config > Shared Prof = No

Location: Ethernet > Connections > *Any Connection Profile*

Ascend-Shared-Profile-Enable

Description: Enables or disables sharing of a RADIUS user file for multiple incoming users.

Note: To apply Shared Profiles on a per RADIUS user profile basis, you have to disable profile sharing on a system-wide basis by setting Ethernet > Mod Config > Shared Prof = No on the MAX

Usage: You can specify one of the following settings:

- Ascend-Shared-Profile-Enable = Shared-Profile-Yes specifies that multiple incoming calls can share this RADIUS user profile.
- Ascend-Shared-Profile-Enable = Shared-Profile-No specifies that multiple incoming calls cannot share a local Connection Profile.
The default value is Shared-Profile-No

Dependencies: For the Ascend-Shared-Profile-Enable attribute to apply, you must disable shared profiles for the MAX as a whole with Ethernet > Mod Config > Shared Prof = No.

Terminal server users can be forced to use unique profiles

The MAX can now force terminal server users to connect using unique profiles.

New parameter

Shared Prof

Description: The MAX can force terminal server users to connect using unique profiles. The Shared Prof parameter in the Ethernet>Mod Config profile or in a Connection profile specifies:

- whether multiple users can share a single Connection profile or a single RADIUS user profile *or*
- whether a single user can have multiple sessions active

This parameter enables multiple incoming calls to share a local Connection profile or a RADIUS users file with Connection profile parameters. Sharing a profile cannot result in two IP addresses sharing the same interface, so this parameter is typically used to share profiles when the caller is assigned an IP address dynamically, which ensures that each caller is assigned a unique address.

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

- Yes means the MAX will allow more than one caller to share the same profile, provided that no IP address conflicts will result.
- No means the MAX will not allow shared profiles.

Note: This feature extends support for the Shared Prof parameter to terminal server users. If Shared Prof is set to No and a user attempts to log in to the MAX terminal server with the same username and password as an already active session, the following message is displayed and the MAX disconnects the user: `***Account Already In Use`

Dependencies: This parameter does not apply to Combinet links or connections that have hard-coded IP addresses. For the Ascend-Shared-Profile-Enable attribute to apply, you must disable shared profiles for the MAX as a whole with Ethernet > Mod Config > Shared Prof = No.

Location: Ethernet > Mod Config, Ethernet > Connections > any profile

See Also: Encaps, Name, Pool # Count, Pool # Start, Recv PW

CR-LF characters act as a single CR in terminal server

The MAX terminal server treats Carriage Return-Line Feed characters as a single Carriage Return.

The MAX terminal server treats a CR-LF pair as a single CR instead of as 2 CRs. A CR-LF pair acts correctly as the end of a line in the terminal server. If a line ends with a CR, the MAX processes it as the end of the line; if it then sees a LF, the MAX ignores it. This change should be transparent to most users because the MAX still recognizes both a CR and a LF as the end of a line.

DNIS support for TCP-CLEAR, X25/PAD, and X25/T3POS calls

MAX units now support Dialed Number Identification Service (DNIS) for TCP-CLEAR-, X25/PAD-, and X25/T3POS-encapsulated calls. DNIS allows the MAX to authenticate incoming callers based on their telephone number.

User interface changes

Previously, if you configured a call with TCP-CLEAR, X25/PAD, and X25/T3POS encapsulation, the Calling # and Called # parameters would not be applicable. Now, you can enter telephone numbers in these fields.

RADVision extensions to V.25bis for DNIS

RADVision requires an extension to the V.25bis implementation. In this release, you can specify the called party number, including the subaddress, to the Incoming Call indication (INC).

User interface changes

The V.25bis Preliminary Draft contains the following syntax for the INC indication:

INC_indication= *phone_number* | INC

Ascend's implementation excluded the *phone_number* argument. This argument specifies number the remote device is calling—the ICLID/ANI (Incoming Call ID/ Automatic Number Identifier) provided by the network. The INC_indication now follows the Preliminary Draft format. The format for *phone_number* argument is:

phone_number [/ *subaddress*]

phone_number is the called party number, and *subaddress* is an optional subaddress.

Enhanced support for T1/PRI conversion

This feature allows the MAX to specify the TypeOfNumber and NumberPlanID fields in the Called Party Number information element when dialing out on a PRI line. Previously, TypeOfNumber was set to Unknown (0) and NumberPlanID was set to ISDN (1) by default and could not be changed.

You can find the details on the context of this feature in the subsection of this release note titled "Enabling the use of a robbed-bit PBX with PRI access lines (PRI-to-T1 Conversion). This section replaces "Enabling the use of a PBX on line 2' in Chapter 2 of the ISP Guides of the MAX products.

User interface changes

This section describes the changes to the user interface.

PRI # Type

Description: T1-PRI:PRI # Type is used for outbound calls made by the MAX on PRI lines so that the switch can properly interpret the phone number dialed. Ask your PRI provider for

details on when to use each of the following settings. This parameter specifies the TypeOfNumber field in the called party's information element.

Note: This parameter applies only to calls placed by devices terminating the inband T1 lines provided by the MAX in a T1-PRI conversion configuration.

Usage: Specify one of the following values:

- National specifies phone numbers within the U.S. (TypeOfNumber=2)
- Intl specifies phone numbers outside the U.S. (TypeOfNumber=1)
- Local specifies phone numbers within your Centrex group. (TypeOfNumber=4)
- Abbrev. specifies that the phone number is abbreviated. (TypeOfNumber=6)
- NetSpecific (currently not implemented) specifies that the network you are connected to understands the phone number (TypeOfNumber=3)
- Unknown (the default) specifies that the phone number is none of the above. (TypeOfNumber=0)

Dependencies: The value you specify for PRI # Type in the Dial Plan profile overrides the value of T1-PRI:PRI # Type in the Line profile if you have enabled the unit's Dial Plan profiles.

Location: Net/T1 > Line Config (Line profile)

See Also: T1-PRI:NumPlanID, NumPlanID (Call and Connection profiles), Modem:NumPlanID (System profile)

NumPlanID

Description: T1-PRI:NumPlanID is used for outbound calls made by the MAX on PRI lines so that the switch can properly interpret the phone number dialed. Ask your PRI provider for details on when to use each of the following settings. This parameter specifies NumberPlanID field in the called party's information element.

Note: This parameter applies only to calls placed by devices terminating the inband T1 lines provided by the MAX in a T1-PRI conversion configuration.

Usage: Specify one of the following values:

- Unknown (NumberPlanID=0)
- ISDN (the default) (NumberPlanID=1)
- Private (NumberPlanID=9)

Dependencies: The value you specify for NumPlanID in the Dial Plan profile overrides the value of T1-PRI:NumPlanID in the Line profile if you have enabled the unit's Dial Plan profiles.

Location: Net/T1 > Line Config (Line profile)

See Also: T1-PRI:PRI # Type, NumPlanID (Call and Connection profiles), Modem:NumPlanID (System profile)

Enabling the use of a robbed-bit PBX with PRI access lines (PRI-to-T1 Conversion)

This feature applies to the MAX 4000/2000/1600 T1 models. Use this section if you have PRI lines from the WAN and need to convert to T1 signaling for support of T1 PBXs. In most cases, you cannot use this feature in combination with digital modems.

The following example configuration uses line 1 to send and receive calls on the WAN and line 2 to handle a PBX for voice service. The MAX emulates a WAN switch, so the PBX on line 2 thinks it is connected to an AT&T or other carrier switch.

Note: The PBX must use 2-state inband with DTMF signaling and must support Senderized (en bloc) digit transmission, because the MAX has a preset time limit on received dialing digits. In addition, the called-party number should be available from the switch (DNIS — Dialed Number Identification Service or called-party information element).

To configure a pair of T1 lines to support a PBX:

- 1 Open Net/T1>Line Config for the second pair of T1 lines on the MAX 4000 (that is, slot 2, or the 20-100 menu).

```
Net/T1
  Line Config
    Name=
    1st Line=Trunk
    2nd Line=Disabled
```

Note: The MAX 2000 has only one pair of T1 lines. These steps apply to the Line profile for lines 1 and 2 in slot 1 (the 10-100 menu).

Note: On the MAX 1600, PRI-to-T1 conversion is available only if the Net/T1 slot card is installed and these steps apply to the Line profile for those lines.

- 2 Set the 2nd Line parameter to Trunk.

```
2nd Line=Trunk
```

- 3 Open the Line 1 subprofile and set the Sig Mode parameter to ISDN.

```
Line 1...
  Sig Mode=ISDN
```

On the MAX 1600, this step applies to line #1 of the Net/T1 slot card.

Note: On the MAX 4000 and 1600, you can also set the first pair of T1 lines (slot 1) for ISDN (PRI) signaling, in which case they become available for outgoing calls from the PBX and can switch incoming calls to the PBX.

- 4 Close the Line 1 subprofile.
- 5 Open the Line 2 subprofile and set the Sig Mode parameter to PBX T1.

```
Line 2...
  Sig Mode=PBX T1
```

On the MAX 1600, this step applies to line #2 of the Net/T1 slot card.

- 6 Set the Rob Ctl parameter as required by the PBX. See the Reference Guide for information on setting this parameter.

```
Line 2...
  Rob Ctl=Wink-Start
```

- 7 Set the T1-PRI:PRI # Type parameter as allowed by your PRI lines provider and appropriate to the calls placed by your PBX places. See the Reference Guide for information on setting this parameter.

Line 2...
T1-PRI:PRI # Type=

- 8** Set the T1-PRI:NumPlanID parameter as required by your PRI lines provider. See the Reference Guide for information on setting this parameter.

Line 2...
T1-PRI:NumPlanID=

- 9** The PBX Type parameters tell the MAX what type of service the PBX expects on its T1 line. In most installations the PBX expects Voice service calls with call progress tones. The value Data does not supply call progress tones or information messages to the user. See the Reference Guide for details.

Line 2...
PBX Type=Voice

- 10** The following two parameters tell the MAX whether or not to convert a call incoming on the PRI line(s) to robbed-bit T1 signaling or to answer the call and perform normal incoming call routing.

Set the Ans Service parameter (Most installations select Voice.)

Line 2...
Ans Service=Voice

NOTE: If you set Ans Svc=Voice, incoming voice service calls on PRI line(s) are converted to T1 signaling on the line outgoing to the PBX. Data service calls are routed according to the MAX's normal incoming call routing, do not go to the PBX and are not converted.

NOTE: If you set Ans Svc=Voice, you cannot configure the MAX for both digital modem operation and PBX-T1 support because all incoming voice service calls are switched to the PBX and none ever reach the digital modems.

- 11** Set the Ans # parameter. Most installations leave this parameter blank. See the Reference Guide for details.

Line 2...
Ans #=

- 12** The following parameters convert the phone number dialed at the PBX to an ISDN PRI format.

Set the Delete Digit parameter. See the Reference Guide for details.

Line 2...
Delete Digit=

- 13** Set the Add Number parameter. See the Reference Guide for details.

Line 2...
Add Number=

- 14** The Call-by-Call parameter adds the appropriate ISDN PRI call setup request for calls dialed out from the PBX.. See the Reference Guide for details.

Line 2...
Call-by-Call=

- 15** Close the Line 2 subprofile.

- 16** Close the T1 profile.

- 17** If you have not already set the Modem:NumPlanID parameter in the System Profile (Sys Config menu), set it now. It determines the numbering plan on outgoing calls. It applies

not only to calls placed by the PBX, but to all outgoing call placed by the MAX. See the Reference Guide for details

Note: NOTE: On MAX models with multiple lines configured for ISDN (that is, PRI), outgoing calls from the PBX use the first available channel on any line configured for ISDN signaling. If you wish to select a PRI line for outgoing calls, the number dialed by the PBX must be prefaced by a dialing prefix set up in the Ch n Trnk Grp Line profile parameter and you must enable trunk groups (by setting the Use Trunk Grps System profile parameter to Yes).

Note: When the MAX forwards an incoming call to the PBX, it does not forward the called-party number.

Enhanced support for outbound modem connections

This feature allows the MAX to specify the TypeOfNumber and NumberPlanID fields in the Called Party Number information element when dialing out on a PRI line. Previously, TypeOfNumber was set to Unknown (0) and NumberPlanID was set to ISDN (1) by default and could not be changed. These two new parameters are used for outbound modem calls.

User interface changes

This section describes the changes to the user interface.

Modem:PRI # Type

Description: Modem:PRI # Type is used for outbound calls made by the MAX on PRI lines so that the switch can properly interpret the phone number dialed. Ask your PRI provider for details on when to use each of the following settings. This parameter specifies the TypeOfNumber field in the called party's information element.

Note: This parameter applies only to calls placed by the digital modems in the MAX; that is, modem dial-out.

Usage: Specify one of the following values:

- National specifies phone numbers within the U.S. (TypeOfNumber=2)
- Intl specifies phone numbers outside the U.S. (TypeOfNumber=1)
- Local specifies phone numbers within your Centrex group. (TypeOfNumber=4)
- Abbrev. specifies that the phone number is abbreviated. (TypeOfNumber=6)
- NetSpecific (currently not implemented) specifies that the network you are connected to understands the phone number (TypeOfNumber=3)
- Unknown (the default) specifies that the phone number is none of the above. (TypeOfNumber=0)

Location: Net/T1 > Line Config (Line profile)

See Also: Modem:NumPlanID, NumPlanID (Call and Connection profiles), T1-PRI:NumPlanID (System profile)

Modem:NumPlanID

Description: Modem:NumPlanID is used for outbound calls made by the MAX on PRI lines so that the switch can properly interpret the phone number dialed. Ask your PRI provider for details on when to use each of the following settings. This parameter specifies NumberPlanID field in the called party's information element.

Note: This parameter applies only to calls placed by the digital modems in the MAX; that is, modem dial-out.

Usage: Specify one of the following values:

- Unknown (NumberPlanID=0)
- ISDN (the default) (NumberPlanID=1)
- Private (NumberPlanID=9)

Location: Net/T1 > Line Config (Line profile)

See Also: Modem:PRI # Type, NumPlanID (Call and Connection profiles), T1-PRI:NumPlanID (System profile)

PRI number plan in outgoing AIM/BONDING and PPP calls

This feature allows the MAX to specify the NumberPlanID field in the Called Party Number information element sent to the switch when dialing out on a PRI line. Previously, NumberPlanID was set to ISDN (1) by default and could not be changed. This feature increased the range of settings available to the PRI # Type parameter (TypeOfNumber field).

User interface changes

This section describes the changes to the user interface. The Pri # Type parameter (Call and Connection profiles) has had additional settings added. The NumPlanIDk parameter (Call and Connection profiles) is new.

L-T/I-P

Description: This parameter has been removed.

U-T/U-P

Description: This parameter has been removed.

PRI # Type

Description: PRI # Type is used for outbound calls made by the MAX on PRI lines so that the switch can properly interpret the phone number dialed. Ask your PRI provider for details

on when to use each of the following settings. This parameter specifies the TypeOfNumber field in the called party's information element.

Usage: Specify one of the following values:

- National (the default) specifies phone numbers within the U.S. (TypeOfNumber=2)
- Intl specifies phone numbers outside the U.S. (TypeOfNumber=1)
- Local specifies phone numbers within your Centrex group. (TypeOfNumber=4)
- Abbrev. specifies that the phone number is abbreviated. (TypeOfNumber=6)
- NetSpecific (currently not implemented) specifies that the network you are connected to understands the phone number (TypeOfNumber=3)
- Unknown specifies that the phone number is none of the above. (TypeOfNumber=0)
- Inherit (Dial Plan profile only) applies to calls placed by a device connected to a local T1 PRI line supplied by a Host/BRI module. If you choose this setting, the caller on the WAN requests the same TypeOfNumber as the caller on the local ISDN BRI line.

Dependencies: The value you specify for PRI # Type in the Dial Plan profile overrides the value of PRI # Type in the Call profile and Connection profile if you have enabled the unit's Dial Plan profiles.

Location: Host/Dual (Host/6)>PortN Menu>Directory (Call profiles), Ethernet>Connections (Connection profiles), System>Dial Plan, Ethernet>Frame Relay, Ethernet>X.25

See Also: NumPlanID, Call-by-Call, T1-PRI:PRI # Type (Line profiles), Modem:PRI# Type (System profile)

NumPlanID

Description: NumPlanID is used for outbound calls made by the MAX on PRI lines so that the switch can properly interpret the phone number dialed. Ask your PRI provider for details on when to use each of the following settings. This parameter specifies NumberPlanID field in the called party's information element.

Usage: Specify one of the following values:

- Unknown (NumberPlanID=0)
- ISDN (the default) (NumberPlanID=1)
- Private (NumberPlanID=9)

Dependencies: The value you specify for NumPlanID in the Dial Plan profile overrides the value of NumPlanID in the Call profile and Connection profile if you have enabled the unit's Dial Plan profiles.

Location: Host/Dual (Host/6)>PortN Menu>Directory (Call profiles), Ethernet>Connections (Connection profiles), System>Dial Plan, Ethernet>Frame Relay, Ethernet>X.25

See Also: PRI # Type, Call-by-Call, T1-PRI:NumPlanID (Line profiles) Modem:NumPlanID (System profile)

PRI number plan in outgoing calls

This feature allows the MAX to specify the NumberPlanID field in the Called Party Number information element sent to the switch when dialing out on a PRI line. Previously, NumberPlanID was set to ISDN (1) by default and could not be changed. This feature makes the MAX compliant with the N1-2 PRI spec.

Overview

When using a local Connection profile or Dial Plan profile to place an outbound PRI call to an NI-2 switch, the MAX was not able to place a 7-digit call as a local call because:

- The called number type could be set in the local profile but not the numbering plan ID. The MAX would always call with the NumberPlanID set to ISDN (1).
- NI-2 spec does not allow calls with “Unknown Number” type and ISDN numbering plan.

As a result, to place a 7-digit call to a NI-2 switch, the MAX had to place it as a 10-digit call with called number type set to “National” and numbering plan set to ISDN Numbering Plan. This feature corrects this problem by making numbering plan selectable in the local Connection and Dial Plan profiles for outbound PRI calls.

User interface changes

This section describes the changes to the user interface.

NumPlanID

Description: NumPlanID is used for outbound calls made by the MAX on PRI lines so that the switch can properly interpret the phone number dialed. Ask your PRI provider for details on when to use each of the following settings. This parameter specifies NumberPlanID field in the called party’s information element.

Usage: Specify one of the following values:

- Unknown (NumberPlanID=0)
- ISDN (the default) (NumberPlanID=1)
- Private (NumberPlanID=9)

Dependencies: The value you specify for NumPlanID in the Dial Plan profile overrides the value of NumPlanID in the Call profile and Connection profile if you have enabled the unit’s Dial Plan profiles.

Location: Host/Dual (Host/6)>PortN Menu>Directory (Call profiles), Ethernet>Connections (Connection profiles), System>Dial Plan, Ethernet>Frame Relay, Ethernet>X.25

See Also: PRI # Type, Call-by-Call, T1-PRI:NumPlanID (Line profiles) Modem:NumPlanID (System profile)

IP routing features

Enable NAT on MAX 1800 and 2000

Ascend NAT (Network Address Translation) functionality is now available on the MAX 1800 and 2000. This functionality makes it possible for the MAX 1800 and 2000 to translate private IP addresses on its local LAN to IP addresses temporarily supplied by a remote access router.

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 MAX 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 MAX 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 MAX 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 that the MAX can access a DHCP server through the remote network.

Single-address NAT and port routing

A MAX 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 MAX.
- By routing packets it receives from the remote network for up to 10 different TCP or UDP ports to specific hosts and ports on the local network.

Note: You can use single-address NAT by setting the Ethernet > NAT > Lan parameter to Single IP Addr.

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

Outgoing connection address translation

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

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

Incoming connection address translation

For incoming calls, the MAX can perform NAT for multiple hosts on the local network using its own IP address. The MAX 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 MAX 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 MAX.

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

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

Translation table size

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

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

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

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

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

Multiple-address NAT

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

The advantages of multiple-address NAT are that hosts on the remote network can connect to specific hosts on the local network, not just specific services such as Web or FTP service, but only if the DHCP server is configured to assign the same address whenever a particular local

host requests an address. Also, network service providers might require multiple-address NAT for networks with more than one host.

When you use multiple-address NAT, hosts on the remote network can connect to any of the official IP addresses that the MAX 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 MAX attempts to perform IP address translation on all packets received. (It cannot distinguish between official and private addresses.)

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

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

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

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

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

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

Configuring single or multiple address NAT

To configure NAT on the MAX:

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

For example:

```
50-C00 NAT
50-C01 NAT...
  >Routing=Yes
  Profile=NATprofile
  Lan=Single IP addr
  FR address=0.0.0.0
  Static Mappings...
  Def Server=N/A
  Reuse last addr=N/A
  Reuse addr timeout=N/A
```

- 2 Enable NAT by setting Routing to Yes. Without this setting, no other setting is valid.
- 3 Set Profile to the name of a Connection profile you want to use NAT.
- 4 FR address refers to Frame Relay. Refer to “NAT for Frame Relay” on page 72 for more information.
- 5 The Static Mappings menu includes 10 Static Mapping *nn* submenus, where *nn* is a value from 01 to 10. Each of these submenus contains parameters for controlling the translation of the private IP addresses to TCP or UDP port numbers when operating in single-address NAT mode. You only need to specify static mappings for connections initiated by devices calling into the private LAN. For sessions initiated by hosts on the private LAN, the MAX generates a mapping dynamically if one does not already exist in the Static Mappings parameters.

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

```
50-C00 NAT
50-C01 NAT...
  Static Mappings...
    Static Mapping 01
      Valid=Yes
      Dst Port #=21
      Protocol=TCP
      Loc Port #=21
      Loc Adrs=181.100.100.102
```

See “Routing incoming sessions to up to 10 servers on the private LAN” on page 73 for information about how to set each parameter.

- 6 Optionally set Def Server to the IP address of a local server to which the MAX routes incoming packets that are *not* routed to a specific server and port. (See “Routing all incoming sessions to the default server” on page 72 for more information.)
- 7 Optionally set Reuse last addr to Yes to continue to use a dynamically assigned IP address. The Reuse addr timeout value specifies the time to use the address. Set it to a number of minutes (up to 1440). Limitations apply, which are described in the *Reference Guide*.
- 8 Exit and save the NAT profile.

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

NAT for Frame Relay

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

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

```
50-C00 NAT
50-C01 NAT...
    Routing=Yes
    Profile=max4
    Lan=Single IP addr
    FR address=0.0.0.0
    Static Mapping...
    Def Server=181.81.8.1
    Reuse last addr=No
    Reuse addr timeout=N/A
```

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

Configuring NAT port routing (Static Mapping submenu)

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

You can configure a NAT port routing

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

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

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

Routing all incoming sessions to the default server

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

- 1 Open the Ethernet > NAT > NAT menu.
- 2 Set the Routing parameter to Yes.
- 3 Set the Profile parameter to the name of an existing Connection profile.
The MAX performs NAT whenever a connection is made with this Connection profile.
The connection can be initiated either by the MAX or by the remote network.
- 4 Set the Lan parameter to Single IP Addr.
- 5 If you previously configured the MAX to route incoming packets for specific TCP or UDP ports (as described in “Routing incoming sessions to up to 10 servers on the private LAN” on page 73).
 - Open each Ethernet > NAT > Static Mapping > Static Mapping *nn* menu (where *nn* is a number between 01 and 10).
 - Set the Valid parameter in each menu to No.
- 6 Set the Def Server parameter to the IP address of the server on the local network to receive all incoming packets from the remote network.
- 7 Press the Esc key to exit the menu.
- 8 Save the changes when prompted.

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

Routing incoming sessions to up to 10 servers on the private LAN

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

- 1 Open the Ethernet > NAT > NAT menu.
- 2 Set the Routing parameter to Yes.
- 3 Set the Profile parameter to the name of an existing Connection profile.
The MAX performs NAT whenever a connection is made with this Connection profile.
The connection can be initiated either by the MAX or by the remote network.
- 4 Set the Lan parameter to Single IP Addr.
- 5 Open the Ethernet > NAT > NAT > Static Mapping menu.
- 6 Open a Static Mapping *nn* menu, where *nn* is a number between 01 and 10.
You use the parameters in each Static Mapping *nn* menu to specify routing for incoming packets sent to a particular TCP or UDP port.
- 7 Set the Valid parameter to Yes.
This enables the port routing specified by the remaining parameters in the menu. Setting this parameter to No disables routing for the specified port.
- 8 Set the Dst Port # parameter to the number of a TCP or UDP port which users outside the private network can access. Each Dst Port # corresponds to a service provided by a server on the local private network. You can use the actual port number as given by the Loc Port # parameter as long as that address is unique for the local private network. See “Well-known ports” on page 74 for information on obtaining port numbers.
The MAX routes incoming packets it receives from the remote network for this port to the local server and port you’re about to specify.

- 9 Set the Protocol parameter to TCP or UDP.
This parameter determines whether the Dst Port # and Loc Port # parameters specify TCP ports or UDP ports.
- 10 Set the Loc Port # to a port corresponding to a service provided by the local servers.
- 11 Set the Loc Adrs parameter to the address of the local server providing the service specified by Loc Port #.
- 12 Exit and save the profile.
Repeat steps 6 through 12 for any additional ports whose packets you want to route to a specific server and port on the local network.
- 13 Open the Ethernet > NAT > NAT menu.
- 14 Set the Def Server parameter to the IP address of a server on the local network that receives any remaining incoming packets from the remote network, that is, any that aren't for ports you've specified in Static Mapping *nn* menus.
- 15 Exit and save the profile.

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

Disabling routing for specific ports

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

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

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

Well-known ports

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

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

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

Parameter Reference

This section describes the new parameters that support this feature.

Routing

Description: The Routing parameter in the NAT profile enables or disables the Network Address functionality of the MAX.

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 MAX 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 MAX 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 MAX 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 that the MAX can access a DHCP server through the remote network.

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

Usage: Select either of the following:

- Yes - Enables NAT
- No - Disables NAT

Dependencies: Route IP in the Ethernet (Mod Config) profile must be set to Yes. NAT automatically turns RIP off, so the address of the MAX is not propagated to the Internet or remote networks.

Dependencies: The Routing parameter in the NAT profile must be set to Yes.

Location: Ethernet > NAT

See Also: Profile, Lan (NAT profile)

Profile

Description: The name of the Connection profile over which the MAX runs network address translation (NAT). Currently, the MAX can specify only one Connection profile to run NAT. 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: Enter a string matching the Name parameter of the Connection profile that runs NAT.

Dependencies: The Routing parameter in the NAT profile must be set to Yes. Route IP in the Connection profile must be set to Yes.

Location: Ethernet > NAT

See Also: Routing (NAT profile)

Lan

Description: Selects whether the MAX is running single-address or multiple-address NAT.

Note: The LAN parameter in the System profile has a different function from the Lan parameter in the NAT profile.

A MAX 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 MAX.
- 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.

Usage: Select either of the following:

- Single IP addr

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

For outgoing calls, the MAX performs NAT on the local network after getting a single IP address from the remote network during PPP negotiation. The MAX does not limit the number of hosts on the local network that can make simultaneous connections to hosts on the remote network. The translations between the local network and the Internet or remote network are dynamic and not preconfigured.

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

- Multiple IP addr

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

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

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

When the first host on the LAN requests access to the remote network, the MAX gets this address through PPP negotiation. When subsequent hosts request access to the remote

network, the MAX asks for an IP address from the DHCP server using a DHCP request packet. The server then sends an address to the MAX from its IP address pool. The MAX uses the dynamic addresses it receives from the server to translate IP addresses on behalf of local hosts. While waiting for an IP address to be offered by the server, corresponding source packets are dropped.

Similarly, for packets received from the WAN, the MAX checks the destination address against its table of translated addresses. If the destination address exists and is active, the MAX forwards the packet. If the destination address does not exist, or is not active, the packet is dropped.

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

Dependencies: The Routing parameter must be set to Yes.

Location: Ethernet > NAT

See Also: Routing (NAT profile)

FR address

Description: The IP address which enables NAT for frame relay connections. Connections using Frame Relay encapsulation can translate local addresses into the single, official address set by this parameter for networking over the wide area network and accessing the Internet.

Usage: Enter an IP address which the remote side of frame relay connections accept as valid. If the remote side of the frame relay connections are routing packets to official IP addresses, then this must be an official IP address.

Dependencies: The Routing parameter in the NAT profile must be set to Yes. The Lan parameter in the NAT profile must be set to Single IP addr. Since this IP address is static, the Reuse last addr parameter does not apply.

Location: Ethernet > NAT

See Also: Routing (NAT profile)

Static Mappings...

Description: A submenu in the NAT profile used to map calls initiated by units outside the local private LAN to servers and services on the local LAN.

Dependencies: The Routing parameter in the NAT profile must be set to Yes. The Lan parameter in the NAT profile must be set to Single IP addr.

Location: Ethernet > NAT

See Also: The parameters in this submenu are Valid, Dst Port #, Protocol, Loc Port # and Loc Adrs.

Valid

Description: This parameter enables or disables the routing of incoming packets for a particular destination port in Static Mappings when the MAX is configured for single-address NAT. This routing is controlled by the parameters in the same Static Mapping nn menu (where nn is a number between 01 and 10).

Usage: Enter Yes or No.

- 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: After you enter a value for this parameter, it does not take effect until the next time the link specified by the Profile parameter is brought up. To make the change immediately, bring the link down and back up.

Dependencies: The Routing parameter in the NAT profile must be set to Yes. The Lan parameter in the NAT profile must be set to Single IP addr.

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

See Also: Routing (NAT profile), Lan (NAT profile), Static Mappings..., Dst Port #, Loc Adrs, Loc Port #, Protocol, Def Server

Dst Port

Description: The number of a TCP or UDP port to which users outside the local private network can send packets to access servers and services on the local LAN. When the MAX is configured for single-address NAT, each Dst Port # corresponds to a service (Loc Port # parameter) provided by a server (Loc Adrs parameter) on the local network; however the actual port number of the service is given by the Loc Port # parameter for which Dst Port # is an alias.

Usage: Enter a port number between 1 and 65535.

Note: After you enter a value for this parameter, it does not take effect until the next time the link specified by the Profile parameter is brought up. To make the change immediately, bring the link down and back up.

Dependencies: The Routing parameter in the NAT profile must be set to Yes. The Lan parameter in the NAT profile must be set to Single IP addr.

- Remember to set the corresponding Loc Port# and Loc Adrs parameters for which this port is an alias.
- The Protocol parameter in the same Static Mapping nn menu determines whether the port you specify is a TCP or UDP port.

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

See Also: Routing (NAT profile), Lan (NAT profile), Static Mappings..

Protocol

Description: When the MAX is configured to perform single-address network address translation (NAT) and to provide services for users outside the private local LAN, 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.

Usage: Enter TCP or UDP.

- 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: After you enter a value for this parameter, it does not take effect until the next time the link specified by the Profile parameter is brought up. To make the change immediately, bring the link down and back up.

Dependencies: The Routing parameter in the NAT profile must be set to Yes. The Lan parameter in the NAT profile must be set to Single IP addr. Valid in Static Mappings must be set to Yes.

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

See Also: Routing (NAT profile), Lan (NAT profile), Static Mappings..., Dst Port #, Loc Adrs, Loc Port #, Valid. Def Server

Loc Port

Description: When the MAX is configured to perform single-address network address translation (NAT) and to provide services for users outside the private local LAN, this parameter specifies the TCP or UDP port of one of the services on the local LAN. The MAX routes packets whose destination port match a setting for Dst Port # to the corresponding Loc Adrs and Loc Port # parameters in the Static Mappings menu.

Usage: Enter the TCP or UDP port number of the local services. The Protocol parameter in the same Static Mapping nn menu determines whether the port you specify is a TCP or UDP port. Settings between 0 and 65535 are allowed. 0 disables static mappings for the corresponding Dst Port #. The default value is 0.

Note: After you enter a value for this parameter, it does not take effect until the next time the link specified by the Profile parameter is brought up. To make the change immediately, bring the link down and back up.

Dependencies: The Routing parameter in the NAT profile must be set to Yes. The Lan parameter in the NAT profile must be set to Single IP addr. Valid in Static Mappings must be set to Yes.

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

See Also: Routing (NAT profile), Lan (NAT profile), Static Mappings..., Dst Port #, Loc Adrs, Protocol, Def Server

Loc Adrs

Description: When the MAX is configured to perform single-address network address translation (NAT) and to provide services for users outside the private local LAN, this parameter specifies the IP address of one of the servers on the local LAN. The MAX routes packets whose destination port match a setting for Dst Port # to the corresponding Loc Adrs and Loc Port # parameters in the Static Mappings menu.

Usage: Enter the IP address of the local server in dotted decimal format. The default value is 0.0.0.0.

Note: After you enter a value for this parameter, it does not take effect until the next time the link specified by the Profile parameter is brought up. To make the change immediately, bring the link down and back up.

Dependencies: The Routing parameter in the NAT profile must be set to Yes. The Lan parameter in the NAT profile must be set to Single IP addr. Valid in Static Mappings must be set to Yes.

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

See Also: Routing (NAT profile), Lan (NAT profile), Static Mappings..., Dst Port #, Loc Port #, Protocol, Def Server

Def Server

Description: The default server in when the MAX is running network address translation (NAT) in single-address mode. The MAX routes incoming packets to the default server when their destination port number does not match an entry in Static Mappings nor does it match a port number dynamically assigned when a local host initiates a TCP / UDP session.

Usage: Enter the IP address of the default server in dotted decimal format. Enter 0.0.0.0 to disable routing of packets to a default server. The default value is 0.0.0.0.

Example: 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 and
- setting the Valid parameter in each of the Static Mapping nn menus to No.

Note: After you enter a value for this parameter, it does not take effect until the next time the link specified by the Profile parameter is brought up. To make the change immediately, bring the link down and back up.

Dependencies: The Routing parameter in the NAT profile must be set to Yes. The Lan parameter in the NAT profile must be set to Single IP addr.

Location: Ethernet > NAT

See Also: Routing (NAT profile), Lan (NAT profile), Static Mappings

Reuse last addr

Description: Specifies that the last IP address given by the remote unit during PPP negotiations should be reused in subsequent PPP negotiations (for the duration specified in the Reuse addr timeout parameter). Reuse last addr applies only to single-address NAT.

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

- Yes - After an IP address is obtained by PPP negotiations, the MAX uses that IP address in all other PPP negotiations as long as the limit set by the Reuse addr timeout parameter has not been exceeded.

Set this parameter to Yes when you need to use the same IP address for TCP applications that might need to reestablish a connection during the session, such as Telnet. For example, suppose a Telnet session is idle for a long period of time and as a result, disconnects, but the Telnet session remains alive. If Reuse last addr = No, when the connection reestablishes, a new IP address is assigned by the PPP negotiations, which creates a problem for Telnet which expected to be using the original IP address.

If the MAX attempts to reuse the IP address and the remote unit rejects the address, it will accept an IP address offered by the remote unit in PPP negotiations.

- No - Each PPP session renegotiates the IP address.

Dependencies: The Routing parameter in the NAT profile must be set to Yes. The Lan parameter in the NAT profile must be set to Single IP addr.

Location: Ethernet > NAT

See Also: Routing (NAT profile), Lan, Reuse addr timeout

Reuse addr timeout

Description: Specifies the time the MAX uses a dynamically assigned IP address when running in single-address NAT mode. See Reuse last addr parameter for details.

Usage: Enter a numeric value from 0 to 1440. The value 0 means the MAX reuses the IP address without a time limit. The maximum setting is 1440 seconds.

Dependencies: The Routing parameter in the NAT profile must be set to Yes. The Lan parameter in the NAT profile must be set to Single IP addr. Reuse last addr parameter must be set to Yes.

Location: Ethernet > NAT

See Also: Routing (NAT profile), Lan, Reuse last addr

NAT Routing

Description: The NAT Routing parameter has the same functionality as the Routing parameter in the NAT profile. This parameter exists for backward compatibility. Do not enable NAT using both NAT Routing and Routing. Only one should be enabled.

Note: NAT has fewer features when enabled from the Mod Config menu. Static mappings and a default server cannot be specified.

Usage: Select either of the following:

- Yes - Enables NAT
- No - Disables NAT

Dependencies: Route IP must be set to Yes. NAT automatically turns RIP off, so the address of the MAX is not propagated to the Internet or remote networks.

Dependencies: The Route IP in the Ethernet (Mod Config) profile must be set to Yes.

Location: Ethernet > Mod Config

See Also: Routing (NAT profile)

NAT Profile

Description: The NAT Profile parameter has the same functionality as the Profile parameter in the NAT profile. This parameter exists for backward compatibility.

Usage: Enter a string matching the Name parameter of the one Connection profile that runs NAT.

Dependencies: The NAT Routing parameter in the NAT profile must be set to Yes. Route IP in the Connection profile must be set to Yes.

Location: Ethernet > Mod Config

See Also: Profile (NAT profile)

NAT Lan

Description: The NAT Lan parameter has the same functionality as the Lan parameter in the NAT profile. This parameter exists for backward compatibility.

Usage: Select either of the following:

- Single IP addr
- Multiple IP addr

Dependencies: The NAT Routing parameter must be set to Yes.

Location: Ethernet > NAT

See Also: Lan (NAT profile)

DHCP services on the MAX 1800 and 2000

A MAX 1800 can now perform a number of Dynamic Host Configuration Protocol (DHCP) services, including responding to DHCP requests to borrow IP addresses, managing Plug and Play requests, and DHCP spoofing.

- 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 MAX. 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 MAX 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 MAX 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 MAX 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 MAX assigns the reserved address.
- If the host is renewing the address it currently has, the MAX 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 MAX always provides the same address.

- If the host is making a new request and there is no IP address reserved for the host, the MAX assigns the next available address from its address pools.

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

Configuring DHCP services

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

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

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

```
20-A00 Mod Config
    DHCP Spoofing...
```

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

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

Always Spoof

Description: Determines how the MAX responds to DHCP requests:

- It can be a Dynamic Host Configuration Protocol (DHCP) server for up to 43 hosts and assign addresses from its own address pools.
- It can perform DHCP spoofing for a single host: it provides a temporary IP address just long enough for a DHCP server on the remote network to provide an actual address. DHCP spoofing is needed only for hosts running APP to authenticate security card users.

When a MAX 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 MAX to be a DHCP server.
- No enables DHCP spoofing. No is the default.

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

Note: Do not enable both DHCP server and BOOTP relay at the same time.

Location: Ethernet > Mod Config > DHCP Spoofing

See Also: DHCP Spoofing, BOOTP Relay Enable

Become Default Router

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

Usage: Press Enter to toggle between choices.

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

Location: Ethernet > Mod Config > DHCP Spoofing

See Also: BOOTP Relay Enable

DHCP PNP Enabled

Description: Determines whether the MAX enables Plug and Play when running in DHCP server mode. In Plug and Play, the MAX assigns an IP address, and returns 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.

Location: Ethernet > Mod Config > DHCP Spoofing

See also: BOOTP Relay Enable

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.

Location: Ethernet > Mod Config > DHCP Spoofing

See Also: Always Spoof

Dial If Link Down

Description: Dial If Link Down applies when both DHCP spoofing and BOOTP relay are enabled. If no wide area network links are active, the MAX performs DHCP spoofing. When set to Yes, as soon as the dialed link is established, the MAX stops DHCP spoofing and acts as a BOOTP relay agent.

Usage: Press Enter to toggle between choices.

- Yes forces the MAX to dial the first Connection profile whenever a it responds to a DHCP client request. Be sure the first Connection profile accesses the DHCP server for which the MAX is spoofing.
 - No lets the MAX 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.

Location: Ethernet > Mod Config > DHCP Spoofing

See Also: BOOTP Relay Enable

Group 1 Count

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

Usage: Enter enter a 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. Press Enter to close the text field.

The default is 1.

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.

Location: Ethernet > Mod Config > DHCP Spoofing

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

Group 2 Count

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

Usage: Enter enter a number between 0 and 20.

If the value is 0, the pool is unavailable.

The default is 0.

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.

Location: Ethernet > Mod Config > DHCP Spoofing

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

Host n IP

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

Usage: Enter the IP address and subnet mask for the host in dotted decimal format.

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.

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 respective Host n Enet parameter. If you disable address assignment by entering 0 . 0 . 0 . 0 / 0 for this parameter, you must set the respective Host n Enet parameter to 000000000000.

Location: Ethernet > Mod Config > DHCP Spoofing

See Also: DHCP Spoofing, Always Spoof

Host n Enet

Description: If the MAX 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 respective Host n IP parameter. When the host sends a DHCP message requesting an IP address, it always receives this address.

Usage: 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.

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 respective Host n IP parameter. If you disable address assignment by entering 000000000000 for this parameter, you must set the respective Host n IP parameter to 0 . 0 . 0 . 0 / 0.

Location: Ethernet > Mod Config > DHCP Spoofing

See Also: DHCP Spoofing, Always Spoof

IP Group 1

Description: The meaning of this parameter depends on whether the MAX 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 MAX 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 MAX 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: Enter an IP address and subnet mask in dotted decimal format.

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.

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.

Location: Ethernet > Mod Config > DHCP Spoofing

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

IP Group 2

Description: If the MAX 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: Enter an IP address and subnet mask in dotted decimal format.

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.

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.

Location: Ethernet > Mod Config > DHCP Spoofing

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

Maximum No Reply Wait

Description: When a MAX 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 MAX waits for a response to the ICMP echo message. If the MAX 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 MAX 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 MAX 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 MAX does not receive the DHCP request during this interval, the MAX stops reserving the address.

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

10 is the default.

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

Location: Ethernet > Mod Config > DHCP Spoofing

See Also: DHCP Spoofing, Always Spoof, Validate IP

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. Renewal Time applies to both DHCP spoofed addresses and DHCP server replies. If the host renews the address before it expires, the MAX provides the same address. Plug and Play addresses always expire in 60 seconds.

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

Location: Ethernet > Mod Config > DHCP Spoofing

Validate IP

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

Location: Ethernet > Mod Config > DHCP Spoofing

See Also: DHCP Spoofing, Always Spoof

Show DHCP command added to terminal-server

The Show DHCP command has been added to the terminal-server interface. It enables you to view DHCP lease and address-assignment information.

Displaying supported DHCP commands

To display the supported DHCP commands, enter the Show DHCP command with a question mark:

```
ascend% show dhcp ?
show dhcp ?           Display help information
show dhcp lease       Display DHCP lease Information
show dhcp address     Display DHCP Address Assignment Information
```

Displaying DHCP leases

To display the users currently assigned addresses via DHCP, enter the Show DHCP Lease command. For example:

```
ascend% show dhcp lease

IP Address      Hardware Address      Netmask           Renew in
10.10.10.2      0080C7300000          255.255.0.0       1130
10.10.10.202    0080C7123450          255.255.0.0       78
12.0.0.101      000078009120          255.255.255.0     12
```

In the output:

- `IP Address` is the address supplied to the client by the MAX, via DHCP.
- `Hardware Address` is the Ethernet (MAC) address of the client configured with the DHCP-supplied address.
- `Netmask` indicates the configured subnet mask for the IP address.
- `Renew in` indicates the number of seconds remaining until the DHCP lease expires for the supplied address. When the lease expires, the client must request another IP address.

Displaying address-assignment information

The MAX displays different output when you enter Show DHCP Address, on the basis of how you have set the Ethernet > Mod Config > DHCP Spoofing > Always Spoof parameter and the Ethernet > DHCP Spoofing > DHCP PNP Enabled parameter.

Displaying addresses when Always Spoof is Yes

Following is sample output the MAX displays when you enter the Show DHCP Address command, and you have set the Always Spoof parameter to Yes and the DHCP PNP Enabled parameter to either Yes or No:

```
ascend% show dhcp address
```

```
DHCP PNP Enabled = Yes
Renewal Time = 1440 seconds
Become Def. Router = Yes
Dial if Link Down = No
Always Spoof = Yes
Validate IP = Yes
Maximum no reply wait = 10 seconds
```

IP Address	Hardware Address	Netmask	In Use
10.10.10.20	00:80:C7:30:00:00	255.255.0.0	Y
10.10.10.21	00:80:C7:12:34:50	255.255.0.0	Y
10.10.10.22	????????????????	255.255.0.0	N
10.10.10.23	????????????????	255.255.0.0	N
10.10.10.24	????????????????	255.255.0.0	N
10.10.10.25	????????????????	255.255.0.0	N
10.10.10.201	????????????????	255.255.0.0	N
10.10.10.202	????????????????	255.255.0.0	N
10.10.10.203	????????????????	255.255.0.0	N
12.178.179.101	00:00:78:00:91:20	255.255.255.0	Y
12.100.123.15	00:80:C7:5B:11:11	255.255.255.0	N
12.100.123.16	00:80:C7:4C:11:11	255.255.255.0	Y

In the output:

- **IP Address** is the address supplied to the host by the MAX, via DHCP.
- **Hardware Address** is the Ethernet (MAC) address of the client configured with the DHCP-supplied address.

Because the MAX learns hardware addresses from either ARP entries or DHCP messages, it is normal for you to see an entry with a hardware address that currently does not have an address assigned to it.

If the MAX requests a hardware address and does not receive a response, the Show DHCP Address command displays ?????????????????? for Hardware Address.
- **Netmask** indicates the configured subnet mask for the IP address.
- **In use** indicates whether or not the IP address is currently assigned to a host.

Displaying addresses when Always Spoof is No and DHCP PNP Enabled is Yes

Following is sample output the MAX displays when you enter the Show DHCP Address command, and you have set the Always Spoof parameter to No and the DHCP PNP Enabled parameter to Yes:

```
ascend% show dhcp address
```

```
DHCP PNP Enabled = Yes
Renewal Time = 1440 seconds
Become Def. Router = Yes
Dial if Link Down = No
Always Spoof = No
Validate IP = Yes
```

IP routing features

Show DHCP command added to terminal-server

Maximum no reply wait = 10 seconds

IP Address	Hardware Address	Netmask	In Use
10.10.10.20	????????????????	255.255.0.0	N

In the output:

- **IP Address** is the address the MAX assigns to PNP clients.
In the example, 10.10.10.20 is derived from the IP address of the MAX unit's Ethernet interface, 10.10.10.19. The displayed IP address is always one greater than that of the MAX.
- **Hardware Address** is the Ethernet (MAC) address of the client configured with the DHCP-supplied address.
Because the MAX learns hardware addresses from either ARP entries or DHCP messages, it is normal for you to see an entry with a hardware address that currently does not have an address assigned to it.
If the MAX requests a hardware address and does not receive a response, the Show DHCP Address command displays ??????????????????.
- **Netmask** indicates the configured subnet mask for the IP address.
- **In use** indicates whether or not the IP address is currently assigned to a host.

Displaying addresses when Always Spoof and DHCP PNP Enabled are No

Following is sample output the MAX displays when you enter the Show DHCP Address command, and you have set the Always Spoof parameter to No and the DHCP PNP Enabled parameter to No:

```
ascend% show dhcp address
```

DHCP PNP Enabled = No
Renewal Time = 1440 seconds
Become Def. Router = Yes
Dial if Link Down = No
Always Spoof = No
Validate IP = Yes
Maximum no reply wait = 10 seconds

IP Address	Hardware Address	Netmask	In Use
10.10.10.17	????????????????	255.255.0.0	N

In the output:

- **IP Address** is address set in the Ethernet > Mod Config > DHCP Spoof > IP Group 1 parameter.
- **Hardware Address** is the Ethernet (MAC) address of the client configured with the DHCP-supplied address.
Because the MAX learns hardware addresses from either ARP entries or DHCP messages, it is normal for you to see an entry with a hardware address that currently does not have an address assigned to it.

If the MAX requests a hardware address and does not receive a response, the Show DHCP Address command displays ??????????????????.

- `Netmask` indicates the configured subnet mask for the IP address.
- `In use` indicates whether or not the IP address is currently assigned to a host.

Changes to the IP router

Changes have been made to Ascend units' IP routing. These include changes to the routing table display and to diagnostic command output.

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

- User interface changes
- New multicast OSPF routes
- Diagnostic mode changes
- Changes to Secure Access Firewall operation

In addition, Ascend units 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						

New multicast OSPF routes

Two new routes have been added to support multicast extensions to OSPF (MOSPF). These are:

224.0.0.5/32	-	local	CP	0	0	0
59593						
224.0.0.6/32	-	local	CP	0	0	0
59593						

Note that 224.0.0.5 and 224.0.0.6 represent the multicast addresses for all OSPF routers on the network and all OSPF designated routers on the network, respectively, and are never forwarded.

On machines which do not support OSPF, these routes are directed instead to the blackhole interface:

224.0.0.5/32	-	bh0	CP	0	0	0
59593						
224.0.0.6/32	-	bh0	CP	0	0	0
59593						

Diagnostic mode changes

The ippacket diagnostic output has been changed.

Modified diagnostic messages

The wording has been changed on these errors:

Table 14. Modified diagnostic messages

Previous message	New message
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

Table 14. Modified diagnostic messages (continued)

Previous message	New message
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 Ascend unit. Previously the packet would be routed by the Ascend unit to the outbound interface, where it would be dropped by the firewall.

Additional information

Keep in mind the following changes to the IP router functionality:

- The MAX now correctly drops certain packets that it would have previously passed. This includes broadcast pings with a source address of 127.0.0.2.
- The MAX did not screen on the source address as defined in RFC1812 4.2.2.11. It does now, except in the case of a source address of 0.0.0.0.

Alphanumeric IP pool names

This feature allows you to assign alphanumeric names to IP address pools.

The Ethernet > Mod Config > WAN Options submenu now contains a Pool Name parameter. You can use this parameter to assign pool names to each of the MAX IP address pools. Certain types of authentication, such as TACACS+, require alphanumeric pool names. When the MAX authenticates a PPP call using TACACS+, it uses the Pool Name to determine which address pool it should use to assign the caller an address. If the Pool Name is not present, or the MAX doesn't find a match, it then attempts to match the Pool Number.

Configuring IP address pool names

To assign a name to an IP address pool:

- 1 Open the Ethernet > Mod Config > WAN Options submenu.
- 2 Set Pool *n* Name to the name of the address pool, where *n* is the number of the address pool.
- 3 Exit and save the Mod Config profile.

Parameter reference

This section describes the new MAX parameter.

Pool Name

Description: Specifies the name of an IP address pool.

Usage: Specify a name. You can enter up to 10 characters. The first character cannot be a number.

Location: Ethernet > Mod Config > WAN Options

Removing routes to a host when the connection is down

In previous releases, the routes associated with Connection profiles were always advertised and added to the MAX 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 MAX remove routes from the routing table when the associated connection is off line, and not advertise these routes.

Overview

The MAX 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 MAX 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:

MAX1 and MAX2 are on the same local LAN.

- MAX1 has a nailed-up line to a remote site. The remote address has a metric of 4.
- MAX2 is a backup connection. It has a remote address with a metric of 7.

Traffic goes through MAX1 because of the lower metric. If MAX1's nailed-up connection goes down, its route to the remote network is still advertised by default. Therefore, the connection specified by MAX2 never comes up.

The new Temporary parameter and Ascend-Temporary-Rtes attribute enable you to specify that the MAX remove the route to an inactive connection's address. The Temporary parameter appears in the IP Options submenu of the Connection profile. The Ascend-Temporary-Rtes attribute appears in a RADIUS user profile. Each one is described in the following section.

User interface changes

Temporary

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

Usage: You can specify one of these settings:

- Yes specifies that the MAX does not advertise the route to a connection when the link is off-line, and removes the route from its routing table, along with all routes dynamically learned on this connection.

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

- No specifies that the MAX advertises 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 MAX unit's routing table, along with all other routes dynamically learned on this connection. All routes age normally. The default value is No.

Example: MAX1 has a nailed connection with an address of 128.50.69.69. MAX1 advertises this route when the connection is up. MAX1 also learns through RIP that the remote side is advertising 198.5.248.72. If the connection goes down and Temporary=Yes, the MAX 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 MAX 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 or a RADIUS user profile. A frame relay link can also have a designated backup Connection profile; if the link goes down, the MAX uses the backup profile for the connection. To specify the backup profile, you use the Backup parameter (on the MAX) or the Ascend-backup attribute (in RADIUS). 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 MAX ignores a setting of Temporary=Yes.
The MAX 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 MAX follows a setting of Temporary=Yes.
The MAX removes routes from the routing table when the frame relay connection goes down.

Ascend-Temporary-Rtes (126)

Description: This attribute specifies whether the MAX stops advertising the route to the address in this RADIUS user profile when the session terminates, and whether the MAX removes this route and all routes dynamically learned on this connection from the routing table.

Usage: You can specify one of these settings:

- Temp-Rtes-No (0) specifies that the MAX advertises the route to the connection found in the Framed-Address attribute, even when the connection is off-line.
The route appears in the MAX unit's routing table, along with all other routes dynamically learned on this connection. The default value is Temp-Rtes-No.
- Temp-Rtes-Yes (1) specifies that the MAX does not advertise the route to a connection when the link is off-line, and removes the route from its routing table, along with all routes dynamically learned on this connection.
The routes are advertised and appear in the routing table only when you re-establish the connection.

Example: MAX1 has a nailed connection with an address of 128.50.69.69. MAX1 advertises this route when the connection is up. MAX1 also learns through RIP that the remote side is advertising 198.5.248.72. If the connection goes down and Ascend-Temporary-Rtes=Temp-Rtes-Yes, the MAX 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 Ascend-Temporary-Rtes=Temp-Rtes-

No, the MAX 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 or a RADIUS user profile. A frame relay link can also have a designated backup Connection profile; if the link goes down, the MAX uses the backup profile for the connection. To specify the backup profile, you use the Backup parameter (on the MAX) or the Ascend-backup attribute (in RADIUS). 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 MAX ignores a setting of Temporary=Yes.
The MAX 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 MAX follows a setting of Temporary=Yes.
The MAX removes routes from the routing table when the frame relay connection goes down.

IP routing display changes

A new “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						

Local DNS host address table option added

To have a fallback when the DNS fails to resolve a hostname successfully, you can now create a local DNS host table that supplies a list of host addresses for important or frequently used connections. You can also specify that the table are automatically updated each time the remote DNS succeeds in resolving a name that is in a list.

Overview

You can now create a local DNS table to provide a list of IP addresses for a specific host name when the remote DNS server fails to resolve the host name. If the local DNS table contains the host name for the attempted connection, it provides the list of IP addresses.

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

IP routing features

Local DNS host address table option added

Automatic updating replaces the existing address list for a host each time the remote DNS server succeeds in resolving a connection to a host that is in the table. You specify how many of the addresses returned by the remote server can be included in the new list.

On the MAX, the table provides includes additional information for each table entry. The information is in the following two fields, which are updated when the system matches the table entry with a host name that was not found by the remote server:

- # Reads (the number of reads since entry was created)
This field is updated each time a local name query match is found in the local DNS table.
- Time of Last Read

Note: The # Reads and Time of last read fields do not appear in the table on a Pipeline 50, Pipeline 75, and Pipeline 130, which do not support SNTP.

You can check the list of host names and IP addresses in the table using the `termserv` command **show dnstab**. Figure 15 shows an example of a DNS table on a MAX. Other terminal server commands show individual entries, with a list of IP addresses for the entry.

Local DNS Table

Name	IP Address	# Reads	Time of last read
1: " "	-----	-----	
2: "server.corp.com."	200.0.0.0	2	Feb 10 10:40:44
3: "boomerang"	221.0.0.0	2	Feb 10 9:13:33
4: " "	-----	-----	
5: " "	-----	-----	
6: " "	-----	-----	
7: " "	-----	-----	

Figure 15. Local DNS table example

New terminal server command changes

New `show` and `dnstab` commands have been added to help you view, edit, or make entries in the DNS table.

show commands

- **show ?** displays a list that includes **dnstab** help.
- **show dnstab** displays the local DNS table.
- **show dnstab ?** displays help for the dnstab editor.

dnstab commands

The terminal server **dnstab** command has three variations:

Table 16. *dnstab commands*

dnstab Command	Description
dnstab	Displays help information about the DNS table.
dnstab show	Displays the local DNS table.
dnstab entry <i>n</i>	<p>Displays a list for entry <i>n</i> in the local DNS table.</p> <p>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.</p> <p>If List Attempt=No, no list is displayed.</p>

Parameter Reference

List Attempt

Description: Enables or disables the DNS List Attempt feature. DNS can return multiple addresses for a hostname in response to a DNS query, but it does not include information about availability of those hosts. Users typically attempt to access the first address in the list. If that host is unavailable, the user must try the next host, and so forth. However, if the access attempt occurs automatically as part of immediate services, the physical connection is torn down when the initial connection fails. To avoid tearing down physical links when a host is unavailable, you can use the List Attempt parameter to enable the MAX to try one entry in the DNS list of hosts, and if that connection fails, to try the next entry, and so on, without losing the WAN session. The List Size parameter specifies the maximum number of hosts listed.

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

- Yes enables a user to try the next host in the DNS list if the first Telnet login attempt fails, which may prevent the physical connection from being torn down.
- No means the connection fails if the first Telnet attempt is refused. For dial-in users, the physical connection is torn down when the initial connection fails.

Dependencies: If List Attempt = No and Enable Local DNS Table = Yes, the local DNS table has only one entry.

Location: Ethernet>Mod Config>DNS

See Also: List Size, Enable Local DNS Table

List Size

Description: Specifies the maximum number of DNS addresses that are made accessible to terminal server sessions in response to a DNS query. List Size also specifies the maximum number of IP address entries in the Local DNS table.

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.

Note: The number of IP addresses displayed with the `dnstab` entry terminal command depends upon the value you set in the List Size parameter.

Usage: Specify a number between 1 and 35. The default is 6.

Example: Following are three possible local DNS table situations:

- You have set List Size=4 and the remote DNS returns 3 addresses, the three addresses replace the entire list of four IP addresses in the local DNS table.
- You have set List Size= 35, and the remote DNS server returns only 4 addresses. The MAX places the four IP addresses in the table and sets the remaining 31 addresses in the list to zero.
- You have just changed the List Size =1. Previously, you had set List Size=10. The next time the table entry for that one IP address is updated, only the first IP address will be retained in the table, all nine others will be set to zero.

Dependencies: This parameter is applicable only when the parameter List Attempt = Yes. A local DNS table is created only if the parameter Enable Local DNS Table= Yes.

Location: Ethernet>Mod Config>DNS

See Also: List Attempt, Enable Local DNS Table

Enable Local DNS Table

Description: Enables the use of a local DNS table that can provide a list of IP addresses for a specific host 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.

Usage: Select Enable Local DNS Table=Yes to enable the local DNS table. No disables the feature. No is the default.

Location: Ethernet Profile: Ethernet > Mod Config > DNS

See Also: The `dnstab` entry terminal command.

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 in the table.

Usage: Loc.DNS Tab Auto Update=Yes to enable automatic updating of the IP addresses in the local DNS table. No disables automatic updating. No is the default.

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.

Dependencies: The Enable Local DNS Table parameter must be set to Yes.

Location: Ethernet Profile: Ethernet > Mod Config > DNS

Configuring the local DNS table

To enable and configure the local DNS table:

- 1 Display Ethernet Profile: Ethernet > Mod Config > DNS menu.
- 2 Select a setting for the List Attempt parameter.
- 3 Specify the list size by setting the List Size parameter.
- 4 Select Enable Local DNS Table=Yes.
The default is No.
- 5 Select a setting for the Loc.DNS Tab Auto Update parameter.

Criteria for valid names in the local DNS table

- Must be unique in the table.
- Must start with an alphabetic character, which may be either upper- or lower-case.
- Must be less than 256 characters
- Names can be local names or fully qualified names that include the domain name.

Periods at the end of names are ignored.

Entering IP addresses in the local DNS table

To enter IP addresses in a local DNS table, you use the DNS table editor from the terminal server. While the editor is in use, the system cannot look up addresses in the table or perform automatic updates. A table *entry* is one of the eight table indexes. It includes the host name, IP address (or addresses), and information fields. To place the initial entries in the table:

- 1 At the terminal server interface, type **dnstab edit**.
Before you make any entries, the table is empty. The editor initially displays zeros for each of the eight entries in the table. To exit the table editor without making an entry, press Enter.
- 2 Type an entry number and press Enter.
A warning appears if you type an invalid entry number. If the entry exists, the current name for that entry appears in the prompt.
- 3 Type the name for the current entry.
If the system accepts the name, it places the name in the table and prompts you for the IP address for the name that you just entered. (For the characteristics of a valid name, see “Criteria for valid names in the local DNS table” on page 105.)

If you enter an invalid name, the system prompts you to enter a valid name.

- 4 Type the IP address for the entry.
If you enter an address in the wrong format, the system prompts you for the correct format. If your format is correct, the system places the address in the table and the editor prompts you for the next entry.
- 5 When you are finished making entries, type the letter **o** and press Enter when the editor prompts you for another entry.

Editing the local DNS table

To edit the DNS table entries, you access the DNS table editor from the terminal server. While the editor is in use, the system cannot look up addresses in the table or perform automatic updates. A table *entry* is one of the eight table indexes. It includes the host name, IP address (or addresses), and information fields. To edit one or more entries in the local DNS table:

- 1 At the terminal server interface, type **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 Enter 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 Replace, accept, or clear the displayed name, as follows:
 - To replace the name, type a new name and press Enter.
 - To accept the current name, press Enter.
 - To clear the name, press the spacebar and then Enter.
If you enter a valid name, the system places it in the table (or leaves it there if you accepted the current name) and prompts you for the corresponding IP address. (For the characteristics of a valid name, see “Criteria for valid names in the local DNS table” on page 105)
If you clear an entry name, all information in all fields for that entry is discarded.
- 4 Either type a new IP address and press Enter, or leave the current address and just press Enter.
 - 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
If the address is in the correct format, the system places it in the table and prompts you for another entry.
- 5 When you are finished making entries, type the letter **o** and press Enter 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 At the terminal server interface, type **dnstab edit** to display the table.
- 2 Type the number of the entry you want to delete and press Return.
- 3 Press the spacebar and then press Return.

UDP Queue Control

You can now control the size of the SNMP and RIP UDP queues. Additional information is reported in the terminal server show udp listen command.

If SNMP or RIP UDP packets arrive at a rate faster than the MAX can process them, then a backlog builds up. This feature lowers the priority of UDP packets destined for the MAX and allows the user to set the size of the SNMP and RIP queues where UDP packets are stored for deferred processing. Prior to this feature, UDP packets destined for the MAX were processed at the same priority as routed packets; they now have a lower priority. Also, prior to this feature, the SNMP UDP queue had no limit and a flood of packets could cause the MAX to run out of memory.

New Parameters

Queue Depth

Description: The maximum number of unprocessed SNMP requests which the MAX saves. If SNMP requests arrive at a rate faster than they can be processed, then a backlog builds up. This parameter sets the maximum depth of the queue. If the queue fills, further packets destined for it are discarded.

Usage: Enter an integer value from 0 to 1024. If you enter 0, the MAX saves SNMP requests until it runs out of memory. 0 is the default.

Note: Setting Queue Depth to 0 is not recommended. An unlimited queue depth could result in an out-of-memory error on the MAX if it receives a flood of packets on its SNMP port.

Location: Ethernet > Mod Config > SNMP options...

See Also: Rip Queue Depth

Rip Queue Depth

Description: The maximum number of unprocessed RIP requests which the MAX saves. If RIP requests arrive at a rate faster than they can be processed, then a backlog builds up. This parameter sets the maximum depth of the queue. If the queue fills, further packets destined for it are discarded. This limit applies to each RIP socket, so if RIP is running on multiple interfaces, this parameter limits the number of requests stored per interface.

Usage: Enter an integer value from 0 to 1024. If you enter 0, the MAX saves RIP requests until it runs out of memory. 50 is the default.

Note: Setting Queue Depth to 0 is not recommended. An unlimited queue depth could result in an out-of-memory error on the MAX if it receives a flood of packets on its RIP port.

Dependencies: This parameter does not apply if the MAX does not listen to RIP updates.

Location: Ethernet > Mod Config > Route Pref...

See Also: Queue Depth, RIP

Displaying UDP statistics and listen

The show udp listen command now shows these additional parameters:

- InQMax - The maximum number of queued UDP packets on the socket (See Queue Depth and Rip Queue Depth parameters.)
- InQLen - The current number of queued packets on the socket
- InQDrops - The number of packets discarded because it would cause InQLen to exceed InQMax
- Total Rx - The total number of packets received on the socket, including InQDrops

An example follows:

```
ascend% show udp listen
udp:
Socket Local Port InQLen InQMax InQDrops Total Rx
0      0      1023    0     1      0      0
1      1      520     0    50      0     532
2      2       7     0    32      0      0
3      3     123     0    32      0      0
4      4    1022     0   128      0      0
5      5     161     0    64      0      0
```

Specifying the metric and preference for offline WAN connections

You can now specify the metric and preference for the MAX to use when a WAN connection is physically down.

User interface changes

The IP Options submenu of the Connection Profile contains two new parameters: DownPreference and DownMetric. The following sections describe each parameter.

DownMetric

Description: This parameter specifies the metric for a route whose associated WAN connection is down.

Usage: Specify an integer. The higher the metric, the less likely that the MAX will use the route. The default metric for online WAN connections is 1. The default metric for offline WAN connections is 7. The metric you specify is in effect only as long as the WAN connection is down.

See Also: DownPreference

DownPreference

Description: This parameter specifies the preference value for a route whose associated WAN connection is down.

Usage: Specify an integer. A higher preference number represents a less desirable route. The default preference for online WAN connections is 60. The default preference for offline WAN connections is 120. The preference you specify is in effect only as long as the WAN connection is down.

Dependencies: Make sure that routes for offline connections have a higher preference number than routes for online connections. The following table lists the factory default values for route preferences.

Table 17. Default route preferences

Route type	Default value
Interface	0
ICMP	30
RIP	100
OSPF ASE	150
OSPF Internal	10
Static	60
Down-Wan	120
Infinite	225

See Also: DownMetric

OSPF features

Specifying an OSPF ASE preference

You can now specify an Autonomous System External (ASE) preference that the MAX uses when it imports OSPF ASEs.

New parameter

In the Ethernet>Mod Config>Route Pref menu, there is now a OSPF ASE Preference parameter. The following section describes this parameter.

OSPF ASE Preference

Description: This specifies the OSPF ASE Preference the MAX uses when importing an ASE.

Usage: Specify a value from 0 to 255. A value of 255 means that the MAX never puts any ASEs into the routing table.

Example: The default route preferences are:

- Connected routes 0
- OSPF internal routes 10
- ICMP routes 30
- Static routes 60
- RIP routes 100
- Unconnected WAN routes 120
- OSPF ASE 150
- Do not use route 255

Dependencies: Keep this additional information in mind.

- When specifying a preference for a route, make sure that routes that are learned from more reliable sources have a lower preference (and are therefore more likely to be used).
- When specifying a preference for a route, you should set a lower preference for connected routes than for disconnected routes.

Location: Ethernet>Mod Config>Route Pref

Specifying the OSPF type of pool advertisement

You can now specify the type of OSPF pool advertisement—either ASE-type1, ASE-type2, or Internal. You can choose any of the settings in the Pool OSPF Adv Type parameter in the Mod Config>WAN options menu.

Support for OSPF NSSAs (RFC 1587)

The MAX now supports OSPF Not So Stubby Areas (NSSAs) as described in RFC 1587. NSSAs allow you to treat complex networks similar to stub areas. This can simplify your network topology and reduce OSPF-related traffic.

Overview

NSSAs are similar to stub areas, except that they allow limited importing of Autonomous System (AS) external routes. NSSAs use type-7 LSAs to import external route information into an NSSA. Type-7 LSAs are similar to type-5 LSAs except that:

- NSSAs can originate and import type-7 LSAs; like stub areas, NSSAs cannot originate or import type-5 LSAs.
- Type-7 LSAs can only be advertised within a single NSSA; they are not flooded throughout the AS as are type-5 LSAs.

When you configure the MAX as an NSSA internal router, you define the type-7 LSAs you want to advertise throughout the NSSA as static routes.

You must also specify whether these type-7 LSAs should be advertised outside the NSSA. If you choose to advertise a type-7 LSA, the NSSA Area Border Router (ABR) converts it to a type-5 LSA, which can then be flooded throughout the AS. If you choose not to advertise a type-7 LSA, it is not advertised beyond the NSSA.

Refer to RFC 1587 for complete information on NSSAs.

Configuring the MAX as an NSSA internal router

Because the MAX cannot be an area border router, when you configure OSPF on the MAX keep in mind that:

- The Area-Type must be the same on all MAX interfaces running OSPF.
- The Area ID (configured in the Area parameter) must be the same on all MAX interfaces running OSPF.

Refer to the documentation that came with your MAX for complete information on configuring OSPF on the MAX.

To configure the MAX as NSSA:

- 1 Select Ethernet > Mod Config > OSPF options.
- 2 Set AreaType to NSSA.
- 3 Exit and save the Mod Config profile.
- 4 Select Ethernet > Static Rtes > *any Static Route profile*.
- 5 Configure a static route to the destination outside the NSSA.
Refer to the documentation that came with your MAX.
- 6 In this static route profile, specify whether you want to advertise this route outside the NSSA:
 - To advertise this route outside the NSSA, set NSSA-Type to Advertise.
 - To not advertise this route outside the NSSA, set NSSA-Type to DoNotAdvertise.
- 7 Exit and save the Static Rtes profile.
- 8 Reset the MAX.

User interface changes

This section describes the changes to the user interface.

Changed parameters

AreaType

The StubNoDefault option has been removed from the AreaType parameter in the Ethernet > Mod Config > OSPF options and the Ethernet > Connection profile > OSPF options submenus.

StubAreaDefaultCost

The StubAreaDefaultCost parameter has been removed from the Ethernet > Mod Config > OSPF options and the Ethernet > Connection profile > OSPF options submenus.

Third Party

The Third Party parameter in the Ethernet > Static Rtes menu is not applicable for NSSAs.

New parameters

NSSA-Type

Description: Specifies whether or not area border routers convert this ASE type-7 to an ASE type-5 LSA. It applies only when the MAX is routing within an OSPF NSSA (that is, where AreaType is set to NSSA on all interfaces running OSPF). ASE type-7s can be imported only from static route definitions. NSSAs are described in RFC 1587.

Usage: Specify one of the following values:

- N/A (the default)
- Advertise (for area border routers to convert this type-7 to a type-5).
- DoNotAdvertise (for area border routers not to convert this type-7 to a type-5)

Dependencies: Keep this additional information in mind:

- Third Party is not applicable when the MAX is configured as an NSSA.
- NSSA-Type is not applicable unless Area-Type is set to NSSA.

Location: Ethernet > Static Rtes > *any Static Rtes profile*

Advertise static routes as OSPF internal LSAs

The MAX can now advertise static routes as internal OSPF Link State Advertisements (LSAs).

This feature allows the MAX to advertise static routes in either internal OSPF Link State Advertisements (LSAs) or Autonomous System External (ASE) LSAs. Previously, the MAX could only advertise static routes as ASE LSAs.

Note that only static routes created in the Static Rtes profile can be advertised as internal LSAs. Static routes created by the MAX from Connection profiles can not be advertised as internal LSAs. Because these dial-up routes are frequently set up and disconnected, advertising them as internal LSAs would cause the SPF tree to be recalculated every time a dial-up call connected or disconnected, substantially degrading performance.

Changed parameter

In the Static Routes profile, the Ase-type parameter is changed to LSA-type and a now includes a new option. The new parameter description is provided below.

LSA-type

Description: This specifies the OSPF ASE type of this link-state advertisement.

Usage: Specify one of the following values:

- ExternalType-1 (the default)
 A type-1 external metric is expressed in the same units as the link-state metric (the same units as interface cost). Type-1 is the default.
- ExternalType-2
 A Type-2 external metric is considered larger than any link state path. Use of type-2 external metrics assumes that routing between autonomous systems is the major cost of routing a packet, and eliminates the need for conversion of external costs to internal link-state metrics.
- Internal
 This indicates that this static route should be advertised in an internal LSA.

Dependencies: Keep this additional information in mind.

- The MAX only advertises the static route if the Static Route gateway has a corresponding entry in a Connection profile.
- When you set LSA-type to Internal, the internal LSA static route appears as a stub area to external OSPF routers.

Location: Ethernet>Connections>OSPF Options, Ethernet>Mod Config>OSPF Option, Ethernet>Static Rtes

See Also: Ospf-Cost

Multicast features

Multicast default route

Previously, when the MAX was operating as a multicast forwarder and concluded that there was no member in a group, it dropped multicast traffic to that group. Now, it forwards that multicast traffic to the Mbone interface instead. This new default behavior is called “multicast default route.”

User interface changes

The only visible change related to this new default is in the output of the Show IGMP Group command. For example, the following output shows that packets are being forwarded to the multicast default route:

```
ascend% show igmp group
IGMP Group address Routing Table   Up Time: 0::16:39:0
Hash      Group Address    Members    Expire time    Counts
N/A       Default route      *(Mbone)   .....        2224862
```

The output contains these fields:

Field	Description
Hash	is an index to a hash table that is displayed for debugging purposes only. N/A means that the default route is not an entry of the hash table.
Group address	is the IP multicast address used. The Default route is the interface on which the multicast router resides.
Members	is the interface ID on which the membership resides. The interface labeled Mbone is the interface on which the multicast router resides.
Expire time	indicates the expiration time of this membership. A string of periods means that the default route never times out.
Counts	indicates the number of packets forwarded.

Configurable multicast group membership timeout

In previous releases, Ascend units configured as a multicast forwarder had a fixed timeout value for IGMP group memberships. The timeout interval is now configurable.

Ethernet>Mod Config>Multicast contains a new parameter named Membership Timeout. This parameter is not applicable unless multicast forwarding is enabled.

When the Ascend unit is configured as a multicast forwarder, it forwards polling messages generated by the multicast router and keeps track of active memberships from its client interfaces. In previous releases, if no client responded to the polling messages within six minutes, the Ascend unit stopped forwarding multicast traffic on that interface. In this release, you can configure the timeout value by specifying a value between 60 seconds and 65535 seconds. The factory default is still six minutes.

To configure the timeout value:

- 1 Open Ethernet>Mod Config>Multicast.
- 2 Configure the Membership Timeout value; for example:

```
Ethernet
  Mod Config...
    Multicast...
      Forwarding=Yes
      Membership Timeout=60
      Mbone Profile=
      Client=Yes
      Rate Limit=100
```

- 3 Close the Ethernet profile.

IPX features

NetWare SAP Home Server Proxy

NetWare SAP Home Server Proxy is a new feature that provides the ability to configure the MAX to forward SAP broadcasts to specified IPX networks to assure that remote users access the same resources as local users.

Background

NetWare SAP Home Server Proxy enables you to give remote users access to the same resources as local users. Rather than relying on the built-in functionality of SAP, the MAX can be configured to direct SAP broadcasts to specified networks.

By default, when you initially load any IPX client software on your PC, a SAP Request packet is broadcast asking for any servers to reply. The first SAP reply received is taken to be the nearest server, and your PC is attached to that server.

If you load your client software from another PC, or use the same PC when travelling, the initial SAP Request could receive responses from different servers and attaching you to different servers. This new feature adds the ability for you to direct SAP Requests to specific networks. The SAP Responses come from servers on these specified networks rather than coming from servers that are near the MAX.

Configuring the NetWare SAP Home Server Proxy

- 1 Open the Ethernet > Connections > Any Connection Profile > IPX Options menu.
- 2 Set the SAP HS Proxy parameter to Yes.
- 3 Specify the IPX network address to which SAP broadcasts will be directed. For example:
`SAP HS Proxy Net#1=CB1123BC`
This indicates any SAP Broadcast Requests received from this user will be directed to IPX network CB1123BC.
- 4 If you want to define other networks, repeat Step 3 for SAP HS Proxy Net#2.

SAP HS Proxy

Description: This parameter specifies whether the MAX performs SAP Home Server Proxy.

Usage: Press Enter to cycle through the choices.

- Yes enables NetWare SAP Home Server Proxy.
- No disables NetWare SAP Home Server Proxy.
No is the default.

Dependencies: The SAP HS Proxy parameter does not apply (SAP HS Proxy=N/A) if IPX routing is disabled (Route IPX=No).

Location: Ethernet > Connections > Any Connection Profile > IPX Options

SAP HS Proxy Net#n (n=1-6)

Description: Specifies an IPX network to which SAP broadcasts should be directed.

Usage: Press Enter to open a text field. Then, type an IPX network number using an 8-digit (4-byte) hexadecimal value. The default is 00000000.

Dependencies: The SAP HS Proxy Net#n parameter does not apply (SAP HS Proxy Net#n=N/A) if either IPX routing is disabled (Route IPX=No) or if SAP Home Server Proxy is disabled (SAP HS Proxy=No).

Location: Ethernet > Connections > Any Connection Profile > IPX Options

SPX spoofing added for IPX

Overview

This feature spoofs the SPX watchdog so that the WAN connection can remain idle while the application(s) requiring it are idle.

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. This feature enables the Ascend unit to allow Netware SPX clients to remain logged in without keeping the WAN connection up in times of inactivity.

To do this, the Ascend unit responds to SPX watchdog requests from the LAN with a "fake" SPX-watchdog-reply packet, and drops any SPX-watchdog-alive packets from the LAN, without sending them on to the WAN.

Note: Routers on both ends of the connection must support this feature for it to function.

New Answer profile option for dial-in NetWare clients

A new Answer parameter has been added to optimize IPX routing for dial-in clients.

In previous releases, the MAX always assumed initially that the far end of an incoming IPX connection was another IPX router. After answering the call, the MAX 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, this default behavior caused problems: the connection would take more than a minute to establish and then the client could not see NetWare servers on the local network. In this release, the Answer profile also contains a Peer parameter to enable the MAX 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 MAX 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 MAX to negotiate inbound IPX

calls as if the far end is a router. The Dialin setting tells the MAX to negotiate inbound IPX calls as if the far end is a dial-in NetWare client.

The following listing 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
```

Dependencies: The MAX must be configured to answer calls for which no configured profile is found (no Connection profile, Names/Passwords profile, or RADIUS entry). The call may require no authentication, or it may 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.

Dependencies: Dial-in NetWare clients do not have their own IPX network, so to enable the MAX to route to dial-in clients, you must specify an IPX Pool number in the Ethernet profile. The network number you specify must be unique within the entire IPX routing domain of the MAX (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 MAX assigns a unique node number to the client as well. It does not send RIP and SAP advertisements across the connection and ignores RIP and SAP advertisements received from the far end. However, it does respond to RIP and SAP queries received from dial-in clients.

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.

Interface changes

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.

RADIUS attribute modified

The Ascend RADIUS attribute 'Ascend-IPX-Route' enables you to configure a static IPX route in a user profile. For example, the traditional argument list for this attribute is:

RADIUS features

ATMP attributes in RADIUS accounting

```
Ascend-IPX-Route = "<profile_name> <network#> [<node#>]  
[<socket#>] [<server_type>] [<hop_count>] [<tick_count>]  
[<server_name>]"
```

Now, you can also add the following:

```
Ascend-IPX-Route = "route-only <network#> <transit_network#>"
```

Note: The ordering of the entries is significant. Route-only entries must come before the definition of the actual intermediate network, within the same dialout profile name (i.e. ipxroute-1).

Table 18. The 5th element

Argument	Description
<network#>	The destination IPX network that you want to define.
<transit_network#>	An intermediate network, that you already know should be used to reach the destination network.

For example, suppose that networks 1 and 2 are connected. Network 1 is connected to a P50. The MAX which is connected to network A has been configured so that it can reach network 1. Without having to define all the information for network 2 again you could use the following in the RADIUS configuration.

```
ipxroute-1 Password = "ascend", User-Service = Dialout-Framed-User
```

```
Ascend-IPX-Route = "route-only 2 1"
```

```
Ascend-IPX-Route = "ipxtest-o 1 000000000001 0451 4 2 12  
COPLAND"
```

For more information, refer to the RADIUS Configuration Guide.

Table 19. The 5th element

Location	Parameters with example values
Ethernet/Mod Config/TServ options.... (Ethernet Profile)	Immediate Modem=Yes Imm. Modem port=5000

RADIUS features

ATMP attributes in RADIUS accounting

For tunneled connections, the RADIUS STOP record previously contained only the Tunneling-Protocol. RADIUS accounting now gathers more information about the tunnel user with the

addition of three new attributes. The new attributes—Ascend-Home-Agent-IP-Addr, Ascend-Home-Agent-UDP-Port, and Ascend-Home-Network-Name—appear in RADIUS Accounting messages when a connection uses ATMP tunneling.

New attributes

Ascend-Home-Agent-IP-Addr

Description: Indicates the IP address of the home agent used for this mobile client.

Example: Following is a RADIUS Accounting record that includes the Ascend-Home-Agent-IP-Addr attribute:

```
Mon Apr 21 02:41:38 1997

User-Name = "JacobP75"
NAS-Identifier = 1.1.1.1
NAS-Port = 10105
Acct-Status-Type = Stop
Acct-Delay-Time = 0
Acct-Session-Id = "111111111"
Acct-Authentic = RADIUS
Acct-Session-Time = 0
Acct-Input-Octets = 215
Acct-Output-Octets = 208
Acct-Input-Packets = 10
Acct-Output-Packets = 10
Ascend-Disconnect-Cause = 1
Ascend-Connect-Progress = 60
Ascend-Data-Rate = 56000
Ascend-PreSession-Time = 1
Ascend-Pre-Input-Octets = 215
Ascend-Pre-Output-Octets = 208
Ascend-Pre-Input-Packets = 10
Ascend-Pre-Output-Packets = 10
Framed-Protocol = PPP
Framed-Address = 2.2.2.2
Tunneling-Protocol = ATMP
Ascend-Home-Agent-IP-Addr = 3.3.3.3
Ascend-Home-Agent-UDP-Port = 5150
Ascend-Home-Network-Name = homenet
```

Dependencies: Accounting-Request packets, generated by the foreign agent, send the Ascend-Home-Agent-IP-Addr attribute at the end of a session, under the following conditions:

- The Accounting-Request packet includes Acct-Status-Type=Stop.
- The session was authenticated and encapsulated by means of Ascend Tunnel Management Protocol (ATMP).

Ascend-Home-Agent-UDP-Port

Description: Identifies the UDP port used when communicating with the home agent.

Example: Following is a RADIUS Accounting record that includes the Ascend-Home-Agent-UDP-Port attribute:

```
Mon Apr 21 02:41:38 1997

User-Name = "JacobP75"
NAS-Identifier = 1.1.1.1
NAS-Port = 10105
Acct-Status-Type = Stop
Acct-Delay-Time = 0
Acct-Session-Id = "111111111"
Acct-Authentic = RADIUS
Acct-Session-Time = 0
Acct-Input-Octets = 215
Acct-Output-Octets = 208
Acct-Input-Packets = 10
Acct-Output-Packets = 10
Ascend-Disconnect-Cause = 1
Ascend-Connect-Progress = 60
Ascend-Data-Rate = 56000
Ascend-PreSession-Time = 1
Ascend-Pre-Input-Octets = 215
Ascend-Pre-Output-Octets = 208
Ascend-Pre-Input-Packets = 10
Ascend-Pre-Output-Packets = 10
Framed-Protocol = PPP
Framed-Address = 2.2.2.2
Tunneling-Protocol = ATMP
Ascend-Home-Agent-IP-Addr = 3.3.3.3
Ascend-Home-Agent-UDP-Port = 5150
Ascend-Home-Network-Name = homenet
```

Dependencies: Accounting-Request packets, generated by the foreign agent, send the Ascend-Home-Agent-UDP-Port attribute at the end of a session, under the following conditions:

- The Accounting-Request packet includes Acct-Status-Type=Stop.
- The session was authenticated and encapsulated by means of Ascend Tunnel Management Protocol (ATMP).

Ascend-Home-Network-Name

Description: Indicates the name of the home network used for this mobile client. This attribute is not present if the home agent is configured in router mode.

Example: Following is a RADIUS Accounting record that displays the Ascend-Home-Network-Name attribute:

```
Mon Apr 21 02:41:38 1997

User-Name = "JacobP75"
NAS-Identifier = 1.1.1.1
NAS-Port = 10105
Acct-Status-Type = Stop
Acct-Delay-Time = 0
Acct-Session-Id = "111111111"
Acct-Authentic = RADIUS
Acct-Session-Time = 0
```



```

Acct-Input-Octets = 215
Acct-Output-Octets = 208
Acct-Input-Packets = 10
Acct-Output-Packets = 10
Ascend-Disconnect-Cause = 1
Ascend-Connect-Progress = 60
Ascend-Data-Rate = 56000
Ascend-PreSession-Time = 1
Ascend-Pre-Input-Octets = 215
Ascend-Pre-Output-Octets = 208
Ascend-Pre-Input-Packets = 10
Ascend-Pre-Output-Packets = 10
Framed-Protocol = PPP
Framed-Address = 2.2.2.2
Tunneling-Protocol = ATMP
Ascend-Home-Agent-IP-Addr = 3.3.3.3
Ascend-Home-Agent-UDP-Port = 5150
Ascend-Home-Network-Name = homenet

```

Dependencies: Accounting-Request packets, generated by the foreign agent, send the Ascend-Home-Network-Name attribute at the end of a session, under the following conditions:

- The Accounting-Request packet includes Acct-Status-Type=Stop.
- The session was authenticated and encapsulated by means of Ascend Tunnel Management Protocol (ATMP).
- The home agent is configured in gateway mode.

RADIUS refers to filter and firewall policies defined in local profile

A new RADIUS attribute enables you to refer to data filters or firewalls defined in local profiles. Previously, you defined the complete filter policy for each user in the appropriate individual RADIUS user profile using the Ascend-Data-Filter attribute. If the filter/firewall policy changed, you had to change the filter definition in each of the affected user's profiles.

You can also use the new RADIUS attribute to specify a firewall policy defined locally on the MAX.

How filtering works with the Filter-Id RADIUS attribute

The RADIUS attribute Filter-Id (11), in the RADIUS user profile, specifies the locally defined data filter or data firewall applied for a user. To assign the same filter or firewall policy to a number of users, you only need to assign the same values to Filter-Id in their RADIUS profiles.

You can specify several filters in a RADIUS user profile, using the Filter-Id attribute in addition to Ascend-Data-Filter and Ascend-Call-Filter. Filter entries apply on a first-match basis. Therefore, the order in which you specify filter entries is significant. The filters and firewalls you specify in the RADIUS user profile are applied for that user the next time the RADIUS user profile is loaded to the MAX.

RADIUS features

RADIUS refers to filter and firewall policies defined in local profile

A match occurs at the first successful comparison between a filter and the packet being examined. When a comparison succeeds, the filtering process stops and the MAX applies the forward or drop action to the packet.

If no comparisons succeed, the packet does not match the filter, and the MAX does *not* forward the packet. When no filter is in use, the MAX forwards all packets. Once you apply a filter to a connection, this default is *reversed*. For security purposes, the MAX does not automatically forward non-matching packets. It requires a rule that explicitly allows those packets to pass.

For more information on filtering, refer to the *Telecommuting and ISP Guide* that came with your MAX unit.

Note: The usage and syntax for Ascend-Data-Filter (and Ascend-Call-Filter) are not modified by Filter-Id.

How firewalls work with the Filter-Id RADIUS attribute

If you specify more than one firewall definition using Filter-Id, only the first firewall definition is applied. If the RADIUS user profile contains a mixture of firewall and filter definitions for Filter-Id, the firewall is applied before any of the filters. The filters are applied after the firewall is applied in the definition sequence described in "How filtering works with the Filter-Id RADIUS attribute."

If you specify a firewall ID for an undefined firewall, a default firewall definition is loaded that allows Telnet packets but not pings.

Ascend-Data-Filter (and Ascend-Call-Filter) do not provide a way to describe a firewall policy. Their usage and syntax are not modified by Filter-Id.

Filter ID numbering

When you create a data filter, you assign it a number between 0 and 199. The number you enter depends on whether you are applying a filter you created using the VT100 interface, or a firewall you created using Secure Access Manager (SAM).

If you are applying a filter created using the VT100 interface, enter the filter number as it appears in the Filters menu.

If you are applying a firewall created with SAM, add 100 to the last 2 digits of the firewall number as it appears in the Firewalls menu. For example, if the number of your firewall is 90-601, enter 101. Refer to your SAM documentation for information on creating firewalls and downloading them to the MAX. The numbering scheme for filters and firewalls is:

- 0 indicates that no filtering is being used (this is the default)
- 1-99 indicates that a filter created using the vt100 interface is being used
- 100-199 indicates that a filter created using SAM is being used.

Configuring Filter-Id in the RADIUS user profile

After you have created a filter on the MAX, you can refer to it in a RADIUS user profile. The following is an example of two data filter profiles and a RADIUS-defined filter applied to a RADIUS user profile

Assume the following two filter profiles are already set up on the MAX are:

```
Filter-id=6
Name=DisallowPing
Out filter 01...Valid=Yes
Out filter 01...Type=IP
Out filter 01...Ip...Forward=No
Out filter 01...Ip...Protocol=6

Filter-id=9
Name=DisallowTelnet
Out filter 01...Valid=Yes
Out filter 01...Type=IP
Out filter 01...Ip...Forward=No
Out filter 01...Ip...Protocol=6
Out filter 01...Ip...Src Port Cmp-Eql
Out filter 01...Ip...Src Port#=23
```

The RADIUS user profile is:

```
someuser Password="ascend"
          User-Service=Framed-User,
          Filter-Id="6",
          Filter-Id="9",
          Ascend-Data-Filter="ip out forward",
          Framed-Protocol=PPP,
          Framed-Address=10.11.1.1,
          Framed-Netmask=255.255.255.0,
          State="p"
```

The first filter is applied, disallowing pings. The second filter disallows Telnet packets. The Ascend-Data-Filter entry allows all IP packets to be forwarded. All pings and Telnet packets will be blocked, but other IP data packets are allowed.

Note: A Telnet directed to another port should be allowed with this configuration.

The following is an example of how Filter-Id can be used to specify a firewall.

- 1 Create a firewall in the SAM program.

The firewall must block all traffic (including Telnets) except ping traffic.

- 2 Download the firewall to the MAX and assign a number, for example,

menu-item 90-101.

- 3 Add the following line to the RADIUS profile in the first example:

Filter-Id="101"

so that the entry reads:

```
someuser Password="ascend"
          User-Service=Framed-User,
          Filter-Id="101",
          Framed-Protocol=PPP,
          Framed-Address=10.11.1.1,
          Framed-Netmask=255.255.255.0,
          State="p"
```

The user should be able to ping into the MAX, but other packets are dropped, since the firewall is applied before the filters are applied.

Filter-Id (11)

Description: This attribute specifies a local data filter or local data firewall profile applied in the current RADIUS user profile. The MAX uses the filter only when it places a call or receives a call using the profile that includes the filter definition. The filters and firewalls specified in the RADIUS user profile are applied for that user the next time the RADIUS user profile is loaded to the MAX.

Usage: You can specify any number of data filters and firewalls. Filter entries apply on a first-match basis, so the order in which you enter the filter entries is significant. If you make changes to a filter in a RADIUS user profile, the changes do not take effect until a call uses that profile.

See Also: Ascend-Data-Filter, Ascend-Call-Filter

NAS-Port-Type added to RADIUS Accounting records

The NAS-Port-Type attribute is displayed in RADIUS Accounting Start, Checkpoint, and Stop records.

Background

To be able to offer different user services, you can configure NAS-Port-Type attribute in a RADIUS Authentication user profile. With this release, NAS-Port-Type is displayed in RADIUS Accounting Start, Checkpoint, and Stop records.

Existing RADIUS attribute

NAS-Port-Type (61)

Description: Specifies the type of service in use for the session.

Some ISPs offer different levels of service on the basis of connection type. To prevent a client from using a capability to which he or she has not subscribed, set the NAS-Port-Type attribute to an appropriate value.

Usage: Specify one of the following settings:

- NAS_Port_Type_Sync specifies a synchronous ISDN connection.
- NAS_Port_Type_Async specifies a call that the MAX routes to a digital modem.
NAS_Port_Type_Async is the default.

Sample Output

Following is a sample RADIUS Accounting Stop record:

```
User-Name = "joe"
NAS-Identifier = 1.1.1.1
NAS-Port = 1091
NAS-Port-Type = Sync    <== (if data call or value = 1)
Acct-Status-Type = Stop
Acct-Delay-Time = 0
Acct-Session-Id = "238536566"
Acct-Authentic = Local
Acct-Session-Time = 28
Acct-Input-Octets = 1050
Acct-Output-Octets = 1060
Acct-Input-Packets = 37
Acct-Output-Packets = 37
Ascend-Disconnect-Cause = 180
Ascend-Connect-Progress = 60
Ascend-Xmit-Rate = 56000
Ascend-Data-Rate = 56000
Ascend-PreSession-Time = 0
Ascend-Pre-Input-Octets = 285
Ascend-Pre-Output-Octets = 266
Ascend-Pre-Input-Packets = 11
Ascend-Pre-Output-Packets = 11
Ascend-First-Dest = 1.1.1.1
Ascend-Multilink-ID = 4
Ascend-Num-In-Multilink = 0
Framed-Protocol = PPP
Framed-Address = 2.2.2.2
```

RADIUS Accounting checkpoint feature

Previously, RADIUS Accounting logged a start and stop record for each user's session. With this new feature, you can configure the MAX to send periodic checkpoint records during a user's session. Should there be a disruption in the network which disconnects active users, you can use these checkpoint records to reconstruct usage in the absence of accounting stop records.

Overview

In a typical environment, RADIUS Accounting allows you to manage complete session information for each of your users. If a disruption in service causes connectivity to be lost before a RADIUS stop record is received for any users, this feature enables you to retrieve information on each user session before the disruption.

By using the new parameter Acct Checkpoint, you can specify the interval at which checkpoint records are sent during each user session.

In the RADIUS detail file, RADIUS Checkpoint records contain the same group of attributes as a RADIUS stop record. However, the value for the Acct-Status-Type attribute in a checkpoint record is the number 3.

Note: When queuing RADIUS Accounting records to be sent to the RADIUS Accounting daemon, the MAX prioritizes start and stop records ahead of checkpoint records.

Acct Checkpoint

Description: Specifies the interval, in minutes, that RADIUS Accounting checkpoint records should be sent for all users.

Usage: Press Enter to open the text field. Type a number from 0 to 60. The default setting is 0, which disables this feature.

Example:

Dependencies: The Acct Checkpoint parameter does not apply (Acct Checkpoint=N/A) if the RADIUS Accounting is not used.

Location: Ethernet > Mod Config > Accounting

NACK option for Local Profiles First parameter

A new option, RNo, has been added to the Local Profiles First parameter. RNo specifies that the MAX first try remote authentication, such as RADIUS, but if a negative acknowledgment (NACK) is received, hangs up and does not try to authenticate against a local profile.

How Local Profiles First authentication works

Two of the options for the Local Profiles First parameter work as they did previously. A new option, RNo, has been added. RNo operates identically to No, except when a NACK comes from the external server. If Local Profiles First is set to No and a NACK is received from the external authentication server, the MAX then check the local Connection and Name/Password profiles. If Local Profiles First is set to RNo and a NACK is received from the external authentication server, the MAX hangs up the call. You might also need to take into account the differences in behavior between remote and local authentication using some authentication methods (see Authentication methods and remote authentication in this note).

Local Profiles First

Description: Local Profiles First enables you to specify whether local authentication is tried before or after remote authentication (RADIUS or other authentication server). A profile is local if it is found in the MAX's VT100 user interface.

Usage: Press Enter to cycle through the choices:

- Local Profiles First = Yes
 - The MAX first checks the local Connection and Name/Password profiles first if the profile exists and password matches then allow connection.
 - if a local profiles matches the caller, except for the password, the MAX checks remote authentication
 - if no local profile matches the caller, the MAX checks remote authentication
- Local Profiles First = No
 - The MAX checks first with the external authentication server (such as RADIUS).

- If the external authentication server doesn't respond in time and there is a TIMEOUT (the server doesn't respond) the MAX checks the local Connection and Name/Password profiles for a match.
- If the external authentication server accepts the request (ACK), the MAX allows the connection.
- If the external authentication server rejects the request (NACK), the MAX checks the local Connection and Name/Password profiles for a match.
- Local Profiles First = RNo
 - The MAX checks first with the external authentication server (such as RADIUS).
 - If the external authentication server doesn't respond in time and there is a TIMEOUT (the server doesn't respond) the MAX checks the Connection and Name/Password profiles for a match.
 - If the external authentication server accepts the request (ACK), the MAX allows the connection.
 - If the external authentication server rejects the request (NACK), the MAX hangs up the connection.

Dependencies: The following authentication methods have problems when Local Profiles First = No or RNo:

- PAP-TOKEN does not produce a challenge if there is a local profile, defeating the purpose of using PAP-TOKEN for authentication.
- PAP-TOKEN-CHAP cannot go beyond one channel in MP+ or MP bundles. The first channel of the bundle is authenticated, but no other channels can be added.
- CACHE-TOKEN for the most part does not work. Although a one-channel session can be initiated, subsequent calls fail until the cached token expires.

Because the remote authentication is tried first in Local Profiles First=No or RNo, the system that has requested external authentication must wait for the remote authentication to time out. This may take longer than the timeout specified for the connection and will cause all connection attempts to fail. To prevent this, set the value for Auth Timeout low enough to not cause the line to be dropped, but still high enough to permit the unit to respond if it is able to. The recommended time is 3 seconds.

Location: Ethernet > Mod Config > Auth

See Also: Auth Timeout

Authentication methods and remote authentication

Some authentication methods do not work the same without a remote authenticator as they do with one. Table 20 shows authentication methods and the specific information you will need to consider if you use a particular method with Local Profiles First=No or =RNo.

Note: You set the values shown in the "Method" column of Table 20 in the Send Auth parameter (Ethernet > Connections > Any Connection Profile > Encaps Options

Table 20. Authentication methods problems

Method	Remote Authentication Considerations if Local Profiles First = No or RNo
PAP	None.
CHAP	None.
PAP-TOKEN	Works either way, but does not produce a challenge if there is a local profile. This defeats the security of using PAP-TOKEN.
PAP-TOKEN-CHAP	Cannot go beyond one channel in MP+ or MP bundles. The first channel of the bundle is authenticated, but no other channels can be added.
CACHE-TOKEN	For the most part, does not work. Although a one-channel session can be initiated, subsequent calls fail until the cached token expires.

SNMP support for Local Profiles First

The SNMP Variable sysAuthPreference now includes the RNo option for Local Profiles First. In the SNMP object that has been added to the Ascend MIB, the following values apply:

Parameter Option	SNMP Equivalent
Local Profiles First = Yes	sysAuthPreference = local-first
Local Profiles First = No	sysAuthPreference = remote-first
Local Profiles First = RNo	sysAuthPreference = remote-no

New SNMP object in Ascend MIB

sysAuthPreference OBJECT-TYPE

SYNTAX INTEGER {
 no-op(1),
 local-first(2),
 remote-first(3),
 remote-no(4)
 }

ACCESS read-write

STATUS mandatory

DESCRIPTION

"An incoming call can be authenticated using a local profile or one from an authentication server such as RADIUS or TACACS.local-first means authenticate from a local profile first, and if that fails, try the authentication server. remote-first means get a profile from the authentication


```
server and authenticate from that and if that fails, try to
authenticate from a local profile. remote-no is similar to
remote-first, except if the external authentication server
NACK the request, then the connection will be denied, i.e. no
search of the local profiles will be made."
::= { systemStatusGroup 10 }
```

Format of RADIUS NAS-Port attribute modified for the MAX

You can now specify that the MAX recognize the same format for the NAS-Port (5) attribute as found on the MAX TNT.

Overview

The NAS-Port (5) specifies the network port on which a call arrives. In past releases of the MAX, the NAS-Port attribute had the format `tllcc`, where

`t`=digital call or analog call

`ll`=line number

`cc`=channel number

By default, the MAX continues to recognize this format. It appears in RADIUS accounting records, and you specify it when restricting a dial-in user to a service, line, and channel. However, you can specify that the MAX recognize the attribute format that specifies a shelf, slot, line, and channel number. This format is used by the MAX TNT.

Changes to the MAX configuration interface

On the MAX, the New NASPort ID parameter enables you to specify the format of the NAS-Port (5) attribute. The following section describes the parameter.

New NASPort ID

Description: Specifies the format the MAX recognizes for the NAS-Port (5) RADIUS attribute.

Usage: Specify one of the following settings:

- Yes specifies that the MAX recognizes the format that specifies a shelf, slot, line, and channel number. This format is the one recognized by the MAX TNT.
- No specifies that the MAX recognizes the five-digit format that specifies the type of service in use, and the line and channel number. The default value is No.

Location: System > Sys Config

Changes to the RADIUS attribute NAS-Port (5)

The new format for NAS-Port is described in the following section.

NAS-Port (5)

Description: Specifies the network port on which the MAX receives a call. The MAX sends NAS-Port to the RADIUS server in an Access-Request packet and an Accounting-Request packet.

Usage: The setting for the NAS-Port attribute is a bit-encoded, zero-based number. To restrict the dial-in user to a slot, line, and channel, enter a value in the following format:

FF SSSS LLLLLL CCCCC

For an ISDN call:

- **FF** specifies the shelf number (always 0 in RADIUS, 1 on the MAX)
- **SSSS** specifies the slot number (0–15)
- **LLLLL** specifies the line number (0–31)
- **CCCCC** specifies the channel number (0–31)

For an analog call, the values are the same, except that the line number can be 0–63, and the channel number is always 1.

Because the value you enter is zero-based, you must add 1 to each component to ascertain the actual slot, line, and channel number. The RADIUS daemon converts the NAS-Port number to decimal on most systems.

Example: To restrict a dial-in user to channel 10 on line 2 for slot 1, specify the following settings:

```
Robin Password="password", NAS-Port=1098
      User-Service=Framed-User,
      Framed-Protocol=PPP,
      Ascend-Assign-IP-Pool=1,
      Ascend-Route-IP=1,
      Ascend-Idle-Limit=300,
      Framed-Routing=None
```

The value NAS-Port=1098 translates to the following NAS port:

- **FF**=shelf 0
- **SSSS**=slot 1
- **LLLLL**=line 2
- **CCCCC**=channel 10

Call-logging feature

Previously, you could configure your MAX to send accounting records to only one RADIUS accounting server. Records contain information that might be useful to other departments. This new feature enables you to send records to another call-logging server.

Call-logging follows the RADIUS Accounting protocol. The call-logging server is a RADIUS accounting server. The MAX communicates with the call-logging server in the same way that it communicates with a RADIUS accounting server.

Overview of call-logging tasks

Call-logging is a way to track information about three types of events:

- Start session. Denotes the beginning of a session with the MAX. Information about this event appears in an logging Start record.
- Stop session. Denotes the end of a session with the MAX. Information about this event appears in a call-logging Stop record.
- Failure-to-start session. Denotes that a login attempt has failed. Information about this event appears in a call-logging Failure-to-start record.

When the MAX recognizes a call-logging event, it sends a call-logging request to the call-logging server. When the call-logging server receives the request, it combines the information into a record and timestamps it. Each type of call-logging record contains attributes associated with an event type, and can show the number of packets the MAX transmits and receives, the protocol in use, the user name and IP address of the client, and so on.

You can use call-logging for either of the following purposes:

- To gather management information. You can use the information in a call-logging record to determine who called, how long the session lasted, and how much traffic occurred during the session.
- To perform troubleshooting. Call-logging records can contain information about how many login failures occurred, and can describe the characteristics of the failed attempts.

Setting up system-wide call-logging values

This section explains how to configure call-logging on a system-wide basis. Some steps are required. Others are optional.

Performing required accounting configuration tasks

When you set up call-logging, you must specify:

- System-wide call-logging parameters
- Call-logging port in `/etc/services`
- Call-logging directory

Specifying system-wide call-logging parameters on the MAX

To set accounting parameters that affect all users on a system-wide basis, perform the following steps at the MAX configuration interface:

- 1 In the External-Auth profile, set Acct-Type =RADIUS.
- 2 Open the Call-logging subprofile.
- 3 For each Host #*n* parameter, specify the IP address of a Call-logging host.
- 4 For the Dst Port parameter, enter the UDP port number you specified, in `/etc/services`, for the authentication process of the daemon. Or, if you used the **incr** keyword with the `-A` option when starting the daemon, add 1 to the number of the UDP port for authentication services and enter the sum.
- 5 For the Key parameter, enter the RADIUS client password, exactly as it appears in the RADIUS clients file.

Specifying the call-logging port

Add to the `/etc/services` file a line identifying the RADIUS daemon's call-logging port. Use the following format:

```
radacct 1646/udp #Call-logging
```

The port number you specify must match the port number indicated by the Dst Port parameter in the Call-logging subprofile.

Specifying the call-logging directory

Create the `/usr/adm/radacct` directory. Or, when starting the daemon, use the `-a` option to specify a different directory in which to store call-logging information. The call-logging process in the daemon creates a file named `detail` in `/usr/adm/radacct`, or in the directory you specify with the `-a` option. The `detail` file contains call-logging records.

Performing optional call-logging configuration tasks

When you configure call-logging, you may optionally specify:

- Timeout value
- Numeric base for the session ID
- Call-logging port

You set each value in the Call-logging subprofile.

Specifying a timeout value

To specify the number of seconds the MAX waits for a response to a call-logging request, set the Acct-Timeout parameter to a value between 1 and 10. The default value is 1.

Specifying the numeric base for the session ID

The Acct-Session-ID attribute is a unique numeric string identified with the session reported in an call-logging packet. The Acct-ID Base parameter controls whether the MAX presents Acct-

Session-ID to the call-logging server in base 10 or base 16. You can specify one of the following settings for the Acct-Id-Base parameter:

- Acct-Base-10 (decimal) indicates that the numeric base is 10. The default value is 10.
- Acct-Base-16 (hexadecimal) indicates that the numeric base is 16.

For example, when you set Acct-Id-Base=Acct-Base-10, the MAX presents a typical session ID to the call-logging server in the following format:

"1234567890"

When you set Acct-Id-Base=Acct-Base-16, the MAX presents the same session ID in the following format:

"499602D2"

Note: Changing the value of Acct-Id-Base while sessions are active creates inconsistencies between the Start and Stop records.

Specifying the call-logging port

To specify the source port the MAX uses to send a call-logging request, set the Dst Port parameter to a value between 0 and 65535. The default value is 0 (zero), which specifies that the Ascend unit can use any port number between 1024 and 2000. You may specify the same source port for authentication and call-logging requests.

Setting up call-logging with dynamic IP addressing

In some networks, the call-logging server requires an IP address for all callers. For callers that receive an IP address from a pool, this requirement presents a problem. During PPP authentication, RADIUS verifies the name and password information, but not the IP address of the caller.

To track calls during the authentication period, you must set up one or more IP address pools, as described elsewhere. Then, in the Rad-Auth-Client subprofile of the External-Auth profile, set Auth-Pool=Yes.

When Auth-Pool=Yes, the MAX includes the caller's assigned IP address as the value of the Framed-Address attribute. The MAX allocates this address from pool #1. (If you do not define pool #1, the call does not have an IP address during authentication.) Because an IP assignment is not usually part of an Access-Request, you must modify the RADIUS daemon.

The assigned IP address might not last the duration of the connection, or it might not be meaningful. Here are five possibilities:

- If Assign-Address=No in the IP-Answer subprofile of the Answer-Defaults profile, and the caller's RADIUS user profile does not supply an IP address for the caller, the MAX returns the IP address to pool #1. However, the address continues to appear in call-logging entries.
- If Assign-Address=No and the caller's RADIUS user profile supplies an IP address for the caller, the MAX returns the IP address to pool #1. The IP address from the user profile appears in call-logging entries.
- If Assign-Address=Yes, and Ascend-Assign-IP-Pool in the RADIUS user profile points to a pool that has no valid IP address, the IP address from pool #1 appears in call-logging entries. The MAX returns the address to the pool only when the call disconnects.

- If Assign-Address=Yes and Must-Accept-Address-Assign=Yes on the MAX, and Ascend-Assign-IP-Pool points to a pool that has a valid IP address, the IP address from that pool appears in call-logging entries for the duration of the call. The MAX returns the address to the pool when the call disconnects.
- If Assign-Address=Yes, Must-Accept-Address-Assign=No, Ascend-Assign-IP-Pool points to a pool that has a valid IP address, and the caller does not specify an address, the IP address from the pool appears in call-logging entries. If the caller does specify an IP address, that address appears in call-logging entries.

Starting the RADIUS daemon with call-logging enabled

To enable call-logging, start the RADIUS daemon with the `-A` option.

When using a flat ASCII file

If you are using a flat ASCII file, enter the following command line:

```
radiusd -A services | incr
```

If you specify the **services** argument, the daemon creates the call-logging process, but only if a line defining the UDP port to use for call-logging appears in the `/etc/services` file. Otherwise, the daemon does not start.

If you specify the **incr** argument, the daemon creates the call-logging process with the UDP port specified as the call-logging port in the `/etc/services` file. If you have not defined the port, the daemon increments the UDP port specified for `radiusd` and uses that port number. This action is the default if you do not specify the `-A` argument.

When using a UNIX DBM database

To start the RADIUS daemon when using a UNIX DBM database, enter the following command line:

```
radiusd.dbm -A services
```

You must specify the **services** argument when you start the daemon in DBM mode.

Understanding call-logging records

This section describes:

- Where call-logging records are stored
- What kinds of packets call-logging uses
- Which attributes appears in each type of packet

Where are call-logging records stored?

The call-logging server writes each record to a log file. If you run an unmodified Ascend RADIUS daemon, the Ascend RADIUS Call-logging file and the Livingston RADIUS Call-logging file have the same name:

```
usr/adm/radacct/host/detail
```

where *host* is the RADIUS client. Because the client of the RADIUS call-logging server is your MAX, *host* is your MAX unit's symbolic host-name, or it's IP address in dotted decimal notation.

What kinds of packets does call-logging use?

Call-logging uses two kinds of packets: Call-logging Start packets and Call-logging Stop packets.

Call-logging Start packets

Call-logging Start packets signal a Start session event. When the MAX begins a terminal-server, bridging, or routing session, and the system authenticates the call or the user logs in, the MAX sends an Call-logging Start packet to the call-logging server. The packet describes the type of session in use and the name of the user opening the session.

The MAX does not send an Call-logging Start packet if a call fails authentication or otherwise fails to log in. In some cases, a session begins with a user login and then authentication follows, such as when a terminal server user chooses PPP or SLIP after login. If User-Service=Login-User, or if User-Service is unspecified, the MAX sends an Call-logging Start packet after login.

Information from an Call-logging Start packet appears in a Start record in the log file.

Call-logging Stop packets

Call-logging Stop packets signal a Stop session or Failure-to-start session event. At the end of a session, including cases in which a user fails authentication, the MAX sends an Call-logging Stop packet. Information from an Call-logging Stop packet appears in a Stop record or Failure-to-start record in the log file.

Non-call-logging attributes in call-logging records

An call-logging record can contain attributes that are not call-logging specific. The following table lists them. Of the attributes listed in the table, only the NAS-Identifier attribute can appear in a Failure-to-start record.

Table 21. Non-call-logging attributes in call-logging records

Attribute	Description
Ascend-Dial-Number (227)	Specifies the phone number of the device that originated the connection.
Caller-Id (31)	Specifies the calling-party number, which is the phone number of the user that has connected to the MAX.
Class (25)	Enables access providers to classify their user sessions. The default value for the Class attribute is null.
Client-Port-DNIS (30)	Specifies the called-party number, which is the phone number the user dials to connect to the MAX.
Framed-Address (8)	Specifies the IP address of the user starting the session. The default value is 0.0.0.0.
Framed-IPX-Network (23)	Specifies the network number of the router at the remote end of the connection. The default value is null.
Framed-Protocol (7)	Specifies the kind of protocol the connection uses. By default, the MAX does not restrict the type of protocol a user can access.
NAS-Identifier (4)	Specifies the IP address of the MAX. This attribute does not appear in an Call-logging-Stop packet for a Failure-start-session event.
NAS-Port (5)	Specifies the network port on which the MAX received the call. This attribute does not appear in an Call-logging-Stop packet for a Failure-start-session event.
User-Name (1)	Specifies the name of the user starting the session.

Call-logging attributes in Start records

The following table lists the call-logging-specific attributes that can appear in a Start record.

Table 22. Call-logging-specific attributes in Start records

Attribute	Description
Acct-Authentic (45)	<p>Specifies the method the MAX used to authenticate an incoming call:</p> <p>RADIUS (1) specifies that RADIUS authenticated the incoming call.</p> <p>Local (2) specifies that the MAX used a local Connection profile, TACACS profile, or TACACS+ profile, or that the MAX accepted the call without authentication.</p>
Acct-Delay-Time (41)	<p>Specifies the number of seconds the MAX has been trying to send the Call-logging packet. In an Call-logging Start packet, this value is 0 (zero).</p>
Acct-Session-Id (44)	<p>Consists of a unique numeric string identified with the bridging, routing, or terminal-server session reported in the Call-logging packet. The string is a random number of up to seven digits.</p> <p>RADIUS correlates the Call-logging Start packet and Call-logging Stop packet with Acct-Session-Id. Its value can range from 1 to 2,137,383,647.</p>
Acct-Status-Type (40)	<p>Requests that have Acct-Status-Type=Start are Call-logging Start packets. The information in these packets appears in Start records.</p> <p>Requests that have Acct-Status-Type=Stop are Call-logging Stop packets. The information in these packets appears in Stop or Failure-to-start records.</p>
Ascend-Session-Svr-Key (151)	<p>Identifies the user session in which a client sends a disconnect or filter-change request to the RADIUS server.</p>
Ascend-User-Acct-Base (142)	<p>Specifies whether the numeric base of the RADIUS Acct-Session-ID attribute is 10 or 16.</p>
Ascend-User-Acct-Host (139)	<p>Specifies the IP address of the RADIUS server to use for the connection.</p>
Ascend-User-Acct-Key (141)	<p>Specifies the RADIUS client password as it appears in the clients file.</p>
Ascend-User-Acct-Port (140)	<p>Specifies a UDP port number for the connection.</p>
Ascend-User-Acct-Time (143)	<p>Specifies the number of seconds the MAX waits for a response to a call-logging request.</p>

Table 22. Call-logging-specific attributes in Start records (continued)

Attribute	Description
Ascend-User-Acct-Type (138)	Specifies the call-logging server(s) to use for the connection.

Call-logging attributes in Stop records

The following table lists the call-logging attributes that can appear in a Stop record.

Table 23. Call-logging-specific attributes in Stop records

Attribute	Description	Conditions for inclusion
Acct-Authentic (45)	Specifies the method the MAX used to authenticate an incoming call: RADIUS (1) specifies that RADIUS authenticated the incoming call. Local (2) specifies that the MAX used a local Connection profile, TACACS profile, or TACACS+ profile, or that the MAX accepted the call without authentication.	Auth-Type parameter not set to RADIUS-Logout. Session must be authenticated.
Acct-Delay-Time (41)	Specifies the number of seconds between the time an event occurred and the time the MAX sent the packet. If RADIUS does not acknowledge the packet, the MAX resends it. The value of Acct-Delay-Time changes to reflect the proper event time.	Auth-Type parameter not set to RADIUS-Logout.
Acct-Input-Octets (42)	Specifies the number of octets the MAX received during the session.	Auth-Type parameter not set to RADIUS-Logout. Session must be authenticated.
Acct-Input-packets (47)	Specifies the number of packets the MAX received during the session.	Auth-Type parameter not set to RADIUS-Logout. Session must be authenticated. A framed protocol must be in use.
Acct-Output-Octets (43)	Specifies the number of octets the MAX sent during the session.	Auth-Type parameter not set to RADIUS-Logout. Session must be authenticated.
Acct-Output-packets (48)	Specifies the number of packets the MAX sent during the session.	Auth-Type parameter not set to RADIUS-Logout. Session must be authenticated. A framed protocol must be in use.

Table 23. Call-logging-specific attributes in Stop records (continued)

Attribute	Description	Conditions for inclusion
Acct-Session-Id (44)	Consists of a unique numeric string identified with the bridging, routing, or terminal-server session reported in the Call-logging packet. The string is a random number of up to seven digits. RADIUS correlates the Call-logging Start packet and Call-logging Stop packet with Acct-Session-Id. Its value can range from 1 to 2,137,383,647.	Auth-Type parameter not set to RADIUS-Logout.
Acct-Session-Time (46)	Specifies the number of seconds the session has been logged in.	Auth-Type parameter not set to RADIUS-Logout. Session must be authenticated.
Acct-Status-Type (40)	Requests that have Acct-Status-Type=Start are Call-logging Start packets. The information in these packets appears in Start records. Requests that have Acct-Status-Type=Stop are Call-logging Stop packets. The information in these packets appears in Stop or Failure-to-start records.	Auth-Type parameter not set to RADIUS-Logout.
Ascend-Connect-Progress (196)	Specifies the state of the connection before it disconnects.	Auth-Type parameter not set to RADIUS-Logout.
Ascend-Data-Rate (197)	Specifies the data rate of the connection in bits per second.	Auth-Type parameter not set to RADIUS-Logout.
Ascend-Disconnect-Cause (195)	Specifies the reason a connection was taken offline.	Auth-Type parameter not set to RADIUS-Logout.
Ascend-Event-Type (150)	Specifies a cold-start notification, informing the call-logging server that the MAX has started up.	For a cold-start notification, the MAX sends values for NAS-Identifier and Ascend-Event-Type in an Ascend-Access-Event-Request packet (code 33). The call-logging server must send back an Ascend-Access-Event-Response packet (code 34), with the correct identifier, to the MAX.
Ascend-First-Dest (189)	Records the destination IP address of the first packet the MAX received on a connection after authentication.	Auth-Type parameter not set to RADIUS-Logout. Session must be authenticated.

Table 23. Call-logging-specific attributes in Stop records (continued)

Attribute	Description	Conditions for inclusion
Ascend-Multilink-ID (187)	Reports the ID number of the Multilink bundle when the session closes.	Auth-Type parameter not set to RADIUS-Logout. Session must be authenticated.
Ascend-Num-In-Multilink (188)	Records the number of sessions remaining in a Multilink bundle when the session closes.	Auth-Type parameter not set to RADIUS-Logout. Session must be authenticated.
Ascend-Number-Sessions (202)	Specifies the number of active user sessions of a given class (as specified by the Class attribute). In the case of multichannel calls, such as MP+ calls, each separate connection counts as a session.	The MAX sends the Ascend-Number-Sessions attribute in Ascend-Access-Event-Request packets. Only RADIUS daemons you customize to recognize packet code 33 respond to these request packets.
Ascend-Pre-Input-Octets (190)	Records the number of octets the MAX received before authentication.	Auth-Type parameter is not set to RADIUS-Logout. The session must be authenticated.
Ascend-Pre-Input-packets (192)	Records the number of packets the MAX received before authentication.	Auth-Type parameter not set to RADIUS-Logout. Session must be authenticated.
Ascend-Pre-Output-Octets (191)	Records the number of octets the MAX sent before authentication.	Auth-Type parameter not set to RADIUS-Logout. Session must be authenticated.
Ascend-Pre-Output-packets (193)	Records the number of packets the MAX sent before authentication.	Auth-Type parameter not set to RADIUS-Logout. Session must be authenticated.
Ascend-PreSession-Time (198)	Specifies the length of time, in seconds, from when a call connected to when it completed authentication.	Auth-Type parameter not set to RADIUS-Logout.
Ascend-User-Acct-Base (142)	Specifies whether the numeric base of the RADIUS Acct-Session-ID attribute is 10 or 16.	None.
Ascend-User-Acct-Host (139)	Specifies the IP address of the RADIUS server to use for the connection.	None.
Ascend-User-Acct-Key (141)	Specifies the RADIUS client password as it appears in the clients file.	None.

Table 23. Call-logging-specific attributes in Stop records (continued)

Attribute	Description	Conditions for inclusion
Ascend-User-Acct-Port (140)	Specifies a UDP port number for the connection.	None.
Ascend-User-Acct-Time (143)	Specifies the number of seconds the MAX waits for a response to a call-logging request.	None.
Ascend-User-Acct-Type (138)	Specifies the call-logging server(s) to use for the connection.	None.

Call-logging attributes in Failure-to-start records

Failure-to-start records can contain only a subset of the information found in Stop records. The following attributes can appear:

- Acct-Delay-Time (41)
- Acct-Session-Id (44)
- Acct-Status-Type (40)
- Ascend-Connect-Progress (196)
- Ascend-Data-Rate (197)
- Ascend-Disconnect-Cause (195)
- Ascend-PreSession-Time (198)

Sample call-logging records

This section provides sample Start and Stop records for the following configurations:

- A Pipeline 25 dialing into a MAX
- A modem calling into a MAX

A Pipeline 25 dialing into a MAX

When a Pipeline 25 dials into a MAX with PPP, the Start record might look like the following:

```
Tue Feb 18 12:00:41 1997 /* Session startup time */
  User-Name="ht-net" /* The name of the Pipeline 25 */
  NAS-Identifier=206.65.212.46 /* The IP address of the MAX */
  NAS-Port=1057 /* Call on channel 2, line 2, slot 2, shelf 1 */
  Acct-Status-Type=Start /* Start record. */
  Acct-Delay-Time=0 /* Always zero for a Start record */
  Acct-Session-Id="1234567" /* Session identification number */
  Acct-Authentic=RADIUS /* RADIUS authentication in use */
  Client-Port-DNIS="3142" /* Called-party number */
  Framed-Protocol=PPP /* PPP call */
  Framed-Address=11.0.0.1 /* IP address of the Pipeline 25 */
```

The Stop record might look like the following:

```
Tue Feb 18 12:02:48 1997 /* Session hangup time */
  User-Name="ht-net" /* The name of the Pipeline 25 */
  NAS-Identifier=206.65.212.46 /* The IP address of the MAX */
  NAS-Port=1057 /* Call on channel 2, line 2, slot 2, shelf 1 */
  Acct-Status-Type=Stop /* Stop record */
  Acct-Delay-Time=18 /* MAX tried to send packet for 18 seconds */
  Acct-Session-Id="1234567" /* Session identification number */
  Acct-Authentic=RADIUS /* RADIUS authentication used */
  Acct-Session-Time=128 /* Number of seconds in session */
  Acct-Input-Octets=2421 /* Bytes received from the Pipeline */
  Acct-Output-Octets=1517 /* Bytes sent to the Pipeline */
  Acct-Input-packets=79 /* Packets received from the Pipeline */
  Acct-Output-packets=47 /* Packets sent to the Pipeline */
  Ascend-Disconnect-Cause=100 /* Session timeout */
  Ascend-Connect-Progress=60 /* LAN session up */
  Ascend-Data-Rate=64000 /* Data rate in bits per second */
  Ascend-PreSession-Time=0 /*Secs from connection to authentication*/
  Ascend-Pre-Input-Octets=174 /* Input octets pre-authentication */
  Ascend-Pre-Output-Octets=204 /* Output octets pre-authentication */
  Ascend-Pre-Input-packets=7 /* Input packets pre-authentication */
  Ascend-Pre-Output-packets=8 /* Output packets pre-authentication */
  Ascend-First-Dest=10.81.44.111 /* Dest IP address of 1st packet */
  Ascend-Multilink-ID=64 /* ID number of Multilink bundle */.
  Ascend-Num-In-Multilink=0 /* # of sessions in Multilink bundle */
  Client-Port-DNIS="3142" /* Called-party number */
  Framed-Protocol=PPP /* PPP call */
  Framed-Address=11.0.0.1 /* IP address of the Pipeline 25 */
```

A modem calling into a MAX

If a modem dials into the MAX to reach its terminal server, the call can only be an unframed call. It cannot be a PPP, MP, or MP+ call. Therefore, the attributes Framed-Protocol and Framed-Address do not appear in the sample records, and Login-Service=Unframed-User.

A Start record might look like the following:

```
Tue Feb 18 12:00:00 1997 /* Session startup time */
  User-Name="Berkeley" /* The name of the modem caller */
  NAS-Identifier=200.65.212.46 /* The IP address of the MAX */
  NAS-Port=1057 /* Call on channel 2, line 2, slot 2, shelf 1 */
  Acct-Status-Type=Start /* Start record. */
  Acct-Delay-Time=0 /* Always zero for a Start record */
  Acct-Session-Id="3456789" /* Session identification number */
  Acct-Authentic=RADIUS /* RADIUS authentication in use */
  Client-Port-DNIS="3143" /* Called-party number */
  Login-Service=Unframed-User /* Modem call */
```

The Stop record might look like the following:

```
Tue Feb 18 12:03:00 1997 /* Session hangup time */
  User-Name="Berkeley" /* The name of the modem caller */
  NAS-Identifier=200.65.212.46 /* The IP address of the MAX */
  NAS-Port=1057 /* Call on channel 2, line 2, slot 2, shelf 1 */
  Acct-Status-Type=Stop /* Stop record */
```

```

Acct-Delay-Time=18 /* MAX tried to send packet for 18 seconds */
Acct-Session-Id="3456789" /* Session identification number */
Acct-Authentic=RADIUS /* RADIUS authentication used */
Acct-Session-Time=128 /* Number of seconds in session */
Acct-Input-Octets=2421 /* Bytes received from the Pipeline */
Acct-Output-Octets=1517 /* Bytes sent to the Pipeline */
Acct-Input-packets=79 /* Packets received from the Pipeline */
Acct-Output-packets=47 /* Packets sent to the Pipeline */
Ascend-Disconnect-Cause=100 /* Session timeout */
Ascend-Connect-Progress=60 /* LAN session up */
Ascend-Data-Rate=64000 /* Data rate in bits per second */
Ascend-PreSession-Time=0 /*Secs from connection to authentication*/
Ascend-Pre-Input-Octets=174 /* Input octets pre-authentication */
Ascend-Pre-Output-Octets=204 /* Output octets pre-authentication */
Ascend-Pre-Input-packets=7 /* Input packets pre-authentication */
Ascend-Pre-Output-packets=8 /* Output packets pre-authentication */
Ascend-First-Dest=10.81.44.111 /* Dest IP address of 1st packet */
Ascend-Multilink-ID=64 /* ID number of Multilink bundle *.
Ascend-Num-In-Multilink=0 /* # of sessions in Multilink bundle */
Client-Port-DNIS="3143" /* Called-party number */
Login-Service=Unframed-User /* Modem call */

```

AppleTalk configurable from RADIUS

You can now set up an AppleTalk connection in a RADIUS user profile and configure static AppleTalk routes in a RADIUS pseudo-user file.

RADIUS attributes added

Three new attributes have been added to the RADIUS dictionary. Two of them enable you to set up an AppleTalk user profile:

- Ascend-Route-Appletalk (118)
- Ascend-Appletalk-Peer-Mode (117)

The third, Ascend-Appletalk-Route, enables you to define a static AppleTalk route.

Ascend-Appletalk-Route (116)

Description: Defines a static AppleTalk route. in a RADIUS pseudo-user profile.

Usage: Create a pseudo-user profile with the first line in the following format:

```
appleroute-num Password="ascend", user-service=Dialout-Framed-User
```

where *num* is a number in a series starting at 1. Then enter one or more static AppleTalk route specifications in the following format:

```
Ascend-Appletalk-Route="net_start net_end zone_name
profile_name"
```


Argument	Description
<i>net_start</i>	The lower limit of the network range for this network. A network range is a range of network numbers set into the port descriptor of the router port and then transmitted through RTMP to the other nodes of the network. Each of the numbers within a network range can represent up to 253 devices. The default is blank.
<i>net_end</i>	The upper limit of the network range for this network. This range defines the networks available for packets routed using the static route. Specify a number between 1 and 65199. If there are other AppleTalk routers on the network, you must configure the network ranges to be identical to the ranges specified on the other routers.
<i>zone_name</i>	The name of the AppleTalk zone associated with this network. A zone is a multicast address containing a subset of the AppleTalk nodes on an internet. Each node belongs to only one zone, but a particular extended network can contain nodes belonging to any number of zones. Zones provide departmental or other groupings of network entities that a user can easily understand. In the Ascend AppleTalk router, zone names are case-insensitive. However, because some routers regard zone names as case-sensitive, the spelling of zone names should be consistent when you configure multiple connections or routers. You can use up to 33 alphanumeric characters. The default is blank.
<i>profile_name</i>	The outgoing RADIUS user profile that the route uses. The default is blank.

Each static route must appear in a user profile. User profile entries for Appletalk static routes are identified by the special name `appleroute-#` and have the following format:

```
appleroute-# Password = "ascend" User-Service = Dialout-Framed-User
    Address 1
    Address 2
    ...
    Address n
```

Address *n* is the actual route associated with this entry.

An example of a static route with the associated connection profiles is:

```
appleroute-1 Password = "ascend" User-Service = Dialout-
Framed-User Ascend-Appletalk-Route = "20 25 testzone1 pipe50"

pipe50 Password = "ascend" User-Service = Dialout-Framed-User,
    User-Service = Framed-User,
    Framed-Protocol = MPP,
    Ascend-Appletalk-Peer-Mode = Appletalk-Peer-Router,
```

```
Ascend-Route-Appletalk = Route-Appletalk-Yes,  
Ascend-Dialout-Allowed = Dialout-Allowed,  
Ascend-Dial-Number = "83272",  
Ascend-Send-Auth = Send-Auth-PAP,  
Ascend-Send-Passwd = "MAX"
```

Dependencies: Ascend-Route-Appletalk must be set to Yes.

See Also: Ascend-Appletalk-Peer-Mode (117), Ascend-Appletalk-Peer-Mode (117)

Ascend-Appletalk-Peer-Mode (117)

Description: Specifies whether the connection is for a single dial-in station or for a router.

Usage: Specify one of the following values:

- Appletalk-Peer-Router (0) specifies that the caller is an AppleTalk router, such as an Ascend Pipeline unit.
- Appletalk-Peer-Dialin specifies that the caller is a dial-in AppleTalk client, such as a single Macintosh dialing in over a modem.

Dependencies: Ascend-Route-Appletalk must be set to Ascend-Route-Appletalk-Yes.

Example: The following example shows a RADIUS user profile for a routed connection:

```
pipe50 Password="pipe50"  
User-Service = Framed-User,  
Framed-Protocol = PPP,  
Ascend-Appletalk-Peer-Mode = Appletalk-Peer-Router,  
Ascend-Route-Appletalk = Route-Appletalk-Yes,  
Ascend-Idle-Limit = 0
```

The following is an example of a RADIUS user profile for a dial-in connection:

```
mac1 Password = "mac1"  
User-Service = Framed-User,  
Framed-Protocol = PPP,  
Ascend-Appletalk-Peer-Mode = Appletalk-Peer-Dialin,  
Ascend-Route-Appletalk = Route-Appletalk-Yes,  
Ascend-Idle-Limit = 0
```

Dependencies: Ascend-Route-Appletalk must be set to Yes.

See Also: Ascend-Appletalk-Peer-Mode (117), Ascend-Appletalk-Route (116)

Ascend-Route-Appletalk (118)

Description: Specifies whether AppleTalk routing is enabled for the connection. When AppleTalk routing is enabled, the connection can forward AppleTalk packets.

Usage: Specify one of the following values:

- Route-Appletalk-No (0) disables AppleTalk routing for this user profile.
- Route-Appletalk-Yes (1) enables AppleTalk routing for this user profile.

The default is No (0).

Dependencies: If you specify Route-Appletalk-Yes, you must set the Ascend-Appletalk-Peer-Mode attribute.

See Also: Ascend-Appletalk-Peer-Mode (117), Ascend-Appletalk-Route (116)

RADIUS Accounting ATMP Notification

Tunneling-Protocol is a new RADIUS Accounting attribute to indicate whether or not a session used the ATMP tunneling protocol.

Tunneling-Protocol

Description: Tunneling-Protocol indicates if a session used the ATMP tunneling protocol.

Example: Below is a sample RADIUS Accounting record which displays the Tunneling-Protocol attribute.

Mon Apr 21 02:41:38 1997

```
User-Name = "JacobP75"
NAS-Identifier = 1.1.1.1
NAS-Port = 10105
Acct-Status-Type = Stop
Acct-Delay-Time = 0
Acct-Session-Id = "111111111"
Acct-Authentic = RADIUS
Acct-Session-Time = 0
Acct-Input-Octets = 215
Acct-Output-Octets = 208
Acct-Input-Packets = 10
Acct-Output-Packets = 10
Ascend-Disconnect-Cause = 1
Ascend-Connect-Progress = 60
Ascend-Data-Rate = 56000
Ascend-PreSession-Time = 1
Ascend-Pre-Input-Octets = 215
Ascend-Pre-Output-Octets = 208
Ascend-Pre-Input-Packets = 10
Ascend-Pre-Output-Packets = 10
Framed-Protocol = PPP
Framed-Address = 2.2.2.2
Tunneling-Protocol = ATMP
```

Dependencies: The Tunneling-Protocol attribute is sent in Accounting-Request packets at the end of a session under the following conditions:

- The Accounting -Request packet has Acct-Status-Stop
- The session was authenticated and encapsulated using the ATMP tunneling protocol.

Specifying preferences for routes configured in

RADIUS

In this release, you can specify a preference for routes you configure in RADIUS.

Changes to RADIUS attributes

This release includes the following new features that enable you to set route preferences in RADIUS:

- A new *preference* option for the existing Framed-Route attribute
- A new Ascend-Preference attribute

Changes to the Framed-Route attribute

You can now specify a route preference for the Framed-Route attribute using the *preference* argument:

```
Framed-Route="host_ipaddr[/subnet_mask] gateway_ipaddr metric [private][name][preference]"
```

The default value is 120. This value is recommended for static IP routes you create in a RADIUS pseudo-user profile. For static routes that you specify using Framed-Route in a dial-in or dial-out RADIUS user profile, we recommend a setting of 60.

New Ascend-Preference attribute

Ascend-Preference (126)

Description: This attribute specifies the preference for a route defined by the Framed-Address attribute in a dial-in or dial-out user profile. Every RADIUS user profile that specifies an explicit IP address using the Framed-Address attribute indicates a static route.

Usage: Specify an integer. The default value is 60. We recommend that you accept this default for dial-in and dial-out user profiles.

Dependencies: Make sure that more desirable routes have a lower preference number. In particular, make sure that routes for connections that are down have a higher preference number than routes for connections that are up. The following table lists the factory default values for route preferences.

Table 24. Default route preferences

Route type	Default value
Interface	0
ICMP	30
RIP	100
OSPF ASE	150

Table 24. Default route preferences (continued)

Route type	Default value
OSPF Internal	10
Static	60
Down-WAN	120
Infinite	225

SecurID authentication for ARA (AppleTalk Remote Access) clients using RADIUS/LOGOUT

Previously, authentication using a SecurID server was available for an ARA client only when you set Auth=RADIUS or Auth=SECURID in the Ethernet>Mod Config>Auth menu. Now, SecurID authentication is available to ARA users when you set Auth=RADIUS/LOGOUT.

How an ARA caller uses SecurID authentication with RADIUS/LOGOUT

An ARA caller can use SecurID authentication in any of the following ways:

- Using a Connection profile
- Using a Password profile
- Using a RADIUS user profile

When Auth=RADIUS/LOGOUT, the user must have the username "SecurID" and no password.

Once the user makes the initial connection, SecurID authentication begins with a pop-up screen on the Macintosh. At this point, the user must enter the "User ID" and "Passcode". If the user enters incorrect values, he or she gets two more tries to authenticate before the connection fails.

If the user is required to enter a new PIN, a pop-up screen prompts for this information. The user has three chances to enter the correct PIN. Once the new PIN is accepted, a pop-up screen instructs the Macintosh user to wait for the token code to change and then to log in with the new PIN and token code.

The SecurID client module must be version 1.3 or later.

For information on setting up a profile to contact an external authentication server, see the *MAX Security Supplement*.

MAX returns to primary RADIUS server after fixed time

Two new parameter Auth Reset Timeout and Acct Reset Timeout have been added to the Ethernet (Mod Config) profile to force the MAX to try to return to the primary RADIUS

authentication and accounting servers. The primary RADIUS servers are set up by the Auth Host #1 and Acct Host #1 parameters.

Parameter Reference

Auth Reset Timeout

Description: This parameter forces the MAX to try to return to the primary RADIUS authentication server; specifically, the server defined by the parameter Auth Host #1.

If a timeout occurs while the MAX was waiting for a reply to an authentication request to the primary RADIUS server; the MAX sends the authentication request to secondary RADIUS server defined by Auth Host #2 and if that fails, Auth Host #3. If either of the secondary servers acknowledges the request, the MAX continues to use that server instead of the primary. Auth Reset Timeout parameter sets the period of time the MAX uses the secondary RADIUS server. At the end of this period of time, the next authentication request the MAX sends to Auth Host #1.

Usage: Enter the period in seconds. Any value from 0 to 86400 is allowed. To disable this feature enter 0 which is equivalent to an infinite number of seconds; that is, the MAX does not return to the primary server as long as the secondary server is replying to requests.

Dependencies: This parameter will be N/A if Auth=None or Auth=TACACS+ in the this profile.

Location: Ethernet Profile: Ethernet > Mod Config > Auth

See Also: Auth Host #1

Acct Reset Timeout

Description: This parameter forces the MAX to try to return to the primary RADIUS accounting server; specifically, the server defined by the parameter Acct Host #1.

If a timeout occurs while the MAX was waiting for a reply to an accounting request to the primary RADIUS server; the MAX sends the accounting request to secondary RADIUS server defined by Acct Host #2 and if that fails, Acct Host #3. If either of the secondary servers acknowledges the request, the MAX continues to use that server instead of the primary. The Acct Reset Timeout parameter sets the period of time the MAX uses the secondary RADIUS server. At the end of this period of time, the next accounting request the MAX sends to Acct Host #1.

Usage: Enter the period in seconds. Any value from 0 to 86400 is allowed. To disable this feature enter 0 which is equivalent to an infinite number of seconds; that is, the MAX does not return to the primary server as long as the secondary server is replying to requests.

Location: Ethernet Profile: Ethernet > Mod Config > Acct

See Also: Acct Host #1

New RADIUS attributes enable call routing to PPTP

You can now use RADIUS to route PPP calls to the PPTP server based on the calling or dialed number, and access more than four PPTP servers. Previously, the MAX did not support routing by PPP-authenticated connection; you had to dedicate a T1 or E1 line for each destination PNS address.

New RADIUS attributes for routing PPTP based on CLID or DNIS

You can now use PPP authentication (CLID and DNIS) to tunnel to PPTP. You are not required to dedicate a T1 line to each destination PNS address, and you are not limited to four PPTP servers, as was the case previously.

Note: It is still possible to dedicate a WAN line to PPTP.

When a PPP call comes in on any WAN line and the authentication process begins, the MAX will first check whether the line is a dedicated PPTP line (the same behavior as previously).

However, if the line is not a PPTP line, the MAX will check the data returned from RADIUS to determine whether:

- CLID or DNIS is supported
- the call is PPTP-based

If the call is a PPTP call and CLID or DNIS is supported, the RADIUS information returned will specify a server endpoint and MAX will route the call through PPTP to the endpoint server. The PPTP server then communicates with the caller.

Configuring PPTP tunnels with CLID or DNIS authentication

You can configure the MAX to route PPP calls authenticated by CLID or DNIS to PPTP by setting the RADIUS parameters described below and in Table 25, which shows the possible values and attribute numbers. "Example RADIUS entries" shows how the attributes appear in a RADIUS entry.

Tunnel-Type (Attribute 64)

Description: Specifies the protocol for the tunnel used with traffic specified by DNIS or CLID.

Usage: Tunnel-Type can have the following values

- PPTP
- L2F (not yet supported)
- L2TP (not yet supported)

Dependencies: Keep this additional information in mind:

- DNIS or CLID must be enabled in the Id Auth parameter of Ethernet Answer profile.
- The MAX must have RADIUS user entries that specify DNIS or CLID.

See Also: Client-Port-DNIS (Attribute 30), used for Called Number authentication.

Tunnel-Medium-Type (Attribute 65)

Description: Specifies the transport medium over which the encapsulated traffic is carried (tunneled).

Usage: Tunnel-Medium-Type can have the following values

- IP
- X25 (not yet supported)
- ATM (not yet supported)

Dependencies: Keep this additional information in mind:

- DNIS or CLID must be enabled in the Id Auth parameter of Ethernet Answer profile.
- The MAX must have RADIUS user entries that specify DNIS or CLID.

See Also: Client-Port-DNIS (Attribute 30), used for Called Number authentication.

Tunnel-Server-Endpoint (Attribute 67)

Description: Specifies the host name or IP address of the destination server.

Usage: Tunnel-Server-Endpoint is a string containing a hostname or IP address for the destination server. If it is a hostname the MAX looks up the host IP address using DNS.

Dependencies: Keep this additional information in mind:

- DNIS or CLID must be enabled in the Id Auth parameter of Ethernet Answer profile.
- The MAX must have RADIUS user entries that specify DNIS or CLID.

See Also: Client-Port-DNIS (Attribute 30), used for Called Number authentication.

Tunnel-Client-Endpoint (Attribute 66)

Description: A string assigned by RADIUS that specifies the name for the unit placing the call. This is used by RADIUS accounting for tracking the session.

Dependencies: Keep this additional information in mind:

- DNIS or CLID must be enabled in the Id Auth parameter of Ethernet Answer profile.
- The MAX must have RADIUS user entries that specify DNIS or CLID.

See Also: Client-Port-DNIS (Attribute 30), used for Called Number authentication.

Tunnel-ID (Attribute 68)

Description: String assigned by RADIUS to each session using CLID or DNIS tunneling. This value is used for accounting when accounting is implemented.

Dependencies: Keep this additional information in mind:

- DNIS or CLID must be enabled in the Id Auth parameter of Ethernet Answer profile.
- The MAX must have RADIUS user entries that specify DNIS or CLID.

See Also: Client-Port-DNIS (Attribute 30), used for Called Number authentication.

Table 25. RADIUS attributes

Attribute	Attribute Values	Attribute Number
Tunnel-Type	1 (Tunnel-Type-PPTP) 2 (Tunnel-Type-L2F) 3 (Tunnel-Type-L2TP)	64
Tunnel-Medium-Type	1 (Tunnel-Medium-Type-IP) 2 (Tunnel-Medium-Type-X25) 3 (Tunnel-Medium-Type-ATM)	65
Tunnel-Client-Endpoint	string (maximum length 253 bytes)	66
Tunnel-Server-Endpoint	string (maximum length 253 bytes)	67
Tunnel-ID	string (maximum length 253 bytes)	68

Example RADIUS entries

The following examples show RADIUS entries for CLID and DNIS. The MAX must have RADIUS user entries that specify DNIS.

CLID RADIUS entry

5105551212 Password = "Ascend-CLID"

Tunnel-Server-Endpoint = "192.168.6.199",

Tunnel-Type = PPTP,

Tunnel-Medium-Type = IP

DNIS RADIUS entry

7894 Password = "Ascend-DNIS"

Tunnel-Server-Endpoint = "eng-lab-199",

Tunnel-Type = PPTP,

Tunnel-Medium-Type = IP

RADIUS now recognizes packet types 33 and 34

In November of 1995, Ascend products supporting RADIUS incorporated two new RADIUS packet types—Ascend-Event-Request (33) and the corresponding Ascend-Event-Response

RADIUS features

RADIUS Accounting-Request packets show whether accounting is enabled or disabled

(34). However, the RADIUS daemon did not recognize these packet types; you had to modify the RADIUS daemon in order to use them. Now, the RADIUS daemon treats the Ascend-Event-Request packet as an Accounting-Request packet, and sends back an Ascend-Event-Response packet.

Changes to RADIUS

An Ascend-Event-Request packet can contain the NAS-Identifier and Ascend-Event-Type attributes. Previously, when the RADIUS server received an Ascend-Event-Request packet, it logged the packet as an unknown request, as in this log entry:

```
Nov 8 16:00:52.352 radiusd[307] Unknown request, code = 33
```

The RADIUS server now treats the Ascend-Event-Request packet as any other accounting request:

```
Nov 8 16:00:52.348 radiusd[307] radrecv: maxhp.1027, id = 9,
```

```
code = 33, length = 32
```

```
NAS-Identifier = 206.65.212.252
```

```
Ascend-Event-Type = Ascend-Coldstart
```

The Ascend-ColdStart value is sent when the MAX boots and is logged to the detail file. The radiusd detail file records this information:

```
Wed Jan 15 14:04:06 1997
```

```
NAS-Identifier = 206.65.212.252
```

```
Ascend-Event-Type = Ascend-ColdStart
```

RADIUS Accounting-Request packets show whether accounting is enabled or disabled

The MAX sends an Accounting-Request packet to indicate when accounting is enabled or disabled. This feature is specified in the RADIUS Accounting RFC 2059.

Overview

RADIUS Accounting-Request packets now enable you to tell whether RADIUS accounting is enabled or disabled on a MAX. Because accounting records can contain time-specific information, you may want visible timestamps indicating when an accounting server begins operating and when it is stopped.

Previously, the Acct-Status-Type attribute, sent in an Accounting-Request packet to the RADIUS accounting server, had two values—Start (1) and Stop (2). The Acct-Status-Type attribute has two new values—7 (Accounting-On) and 8 (Accounting-Off).

Accounting-Request packet when RADIUS accounting is enabled

The MAX sends an Accounting-Request packet with Acct-Status-Type attribute set to 7, Accounting-On, under the following conditions:

- You boot the MAX and Acct=RADIUS in the Ethernet > Mod Config > Accounting menu.
- You set the Acct parameter RADIUS and save the configuration while the MAX is running.

The MAX retransmits the request until it receives an Accounting-Response packet from the RADIUS accounting server, or until the MAX reaches a limit of ten retries for each accounting server it attempts to reach.

The Accounting-Request packet contains these attributes and values:

- NAS-Identifier (4) with the IP address of the MAX
- Acct-Status-Type (40) with the value 7
- Acct-Delay-Time(41) with the number of seconds the MAX has been trying to send the packet without receiving an acknowledgement from the accounting server.

Accounting-Request packet when RADIUS accounting is disabled

The MAX sends an Accounting-Request packet is sent with Acct-Status-Type set to 8, Accounting-Off, under the following conditions:

- You reset the MAX.
- You set Acct to either None or TACACS+ and save the configuration while the MAX is running.
- You change the setting of the Auth parameter in the Ethernet>Mod Config>Auth menu from RADIUS to RADIUS/LOGOUT.

When Auth=RADIUS/LOGOUT, Acct is set to N/A, and RADIUS accounting is disabled.

Except when the MAX is being reset, it retransmits the request until it receives an Accounting-Response packet from RADIUS, or until the MAX reaches a limit of ten retries for each accounting server it attempts to reach. When the MAX is being reset, it sends only one Accounting-Request packet, and does not retransmit the request. The MAX does not send an Accounting-Request packet in response to a power failure.

The Accounting-Request packet contains these attributes and values:

- NAS-Identifier (4) with the IP address of the MAX
- Acct-Status-Type (40) with the value 8
- Acct-Delay-Time(41) with the number of seconds the MAX has been trying to send the packet without receiving an acknowledgement from the accounting server.

Configuring Data Over Voice Bearer Service (DOVBS)

in RADIUS

In previous releases, you were limited to specifying four phone numbers that the Ascend unit would handle as data calls, even if they came in as voice calls. These numbers are known as DNIS (Dial Number Information Service) numbers. For some purposes, the limitation of four numbers is restrictive. Therefore, in this release, you can use RADIUS to specify up to 100 phone numbers.

Configuring DNIS numbers in RADIUS

This feature addresses those locations in North America that charge significantly higher rates for digital bearer service than voice bearer service. In addition, this feature can be used by larger ISPs that resell services to smaller ISPs, identifying them by DNIS number. To configure DNIS numbers in RADIUS, follow these steps:

- 1 Create the first line of a pseudo-user entry using the User-Name and Password attributes. You create a pseudo-user to store information that the Ascend unit can query—in this case, in order to store DNIS numbers. You can configure pseudo-users for both global and Ascend unit-specific configuration control of DNIS numbers. The Ascend unit adds the unit-specific information in addition to the global information.

For a unit-specific DNIS configuration, specify the first line of a pseudo-user entry in this format:

```
dovbs-<unit_name>-<num> Password="ascend"
```

For a global DNIS configuration, specify the first line of a pseudo-user entry in this format:

```
dovbs-<num> Password="ascend"
```

<unit_name> is the system name of the Ascend unit—that is, the name specified by the Name parameter in the System Profile. <num> is a number in a sequential series, starting at 1.

- 2 For each pseudo-user entry, specify one or more DNIS numbers using the Client-Port-DNIS attribute.

Specify each DNIS number in this format:

```
Client-Port-DNIS="<DNIS number>"
```

Client-Port-DNIS specifies the called-party number, indicating the phone number dialed by the user to connect to the Ascend unit. You can specify up to 100 DNIS numbers; if you specify more than 100, the remaining numbers are ignored.

Consider this example:

Five small ISPs connect to a larger ISP using DOVBS with the following numbers:

- 4165551111
- 4165552222
- 4165553333
- 4165554444
- 4165555555

The pseudo-user profile for an Ascend unit named "Toronto" looks like this one:

```
dovbs-Toronto-1 Password = "ascend"
```

```
Client-Port-DNIS = "5105551111"  
Client-Port-DNIS = "5105552222"  
Client-Port-DNIS = "5105553333"  
Client-Port-DNIS = "5105554444"  
Client-Port-DNIS = "5105555555"
```

How the Ascend unit learns about DNIS entries

When you have properly configured the pseudo-user profile, RADIUS loads DNIS numbers whenever you power on or reset the Ascend unit, when you select the Upd Rem Cfg command from the Sys Diag menu, or when you use an update command in SNMP. RADIUS loads the numbers in this way:

- 1 RADIUS looks for entries having the format dovbs-<unit_name>-1, where <unit_name> is the system name.
- 2 If at least one entry exists, RADIUS loads all existing entries with the format dovbs-<unit_name>-<num>. The variable <num> is a number in a sequential series, starting with 1.
- 3 The Ascend unit queries dovbs-<unit_name>-1, then dovbs-<unit_name>-2, and so on, until it receives an authentication reject from RADIUS.
- 4 Once the host-specific numbers are loaded, RADIUS loads the global configuration entries; these configurations have the format dovbs-<num>.
- 5 The Ascend unit queries dovbs-1, then dovbs-2, and so on, until it receives an authentication reject from RADIUS.
- 6 The Ascend unit checks the DNIS number of each incoming call against the phone numbers it has loaded.

New parameter for generating a second accounting Start record

In this release, when a user logs in using the terminal server and starts a PPP session, the Ascend unit can cause RADIUS to generate a second accounting Start record containing the user's Framed-Address and Framed-Protocol attributes.

User interface changes

Framed Addr Start

Description: This parameter specifies whether the Ascend unit sends a second accounting Start record to the RADIUS server when the Framed-Address and Framed-Protocol attributes are assigned to a user transferring to a framed protocol (such as PPP or SLIP).

Usage: You can specify one of these settings:

- Yes indicates that the Ascend unit sends a second accounting Start record.
- No indicates that the Ascend unit does not send a second accounting Start record.

No is the default.

Location: Ethernet profile: Ethernet>Mod Config>Auth

New RADIUS and terminal server support for Point to Point Tunneling Protocol (PPTP)

You can now select a PPTP tunnel on a per-user basis, route a terminal-server session to a PPTP server, and specify either a host name or an IP address in the route to a PPTP server.

Overview

In past releases, you could not select a PPTP tunnel on a per-user basis. In this release, you can assign each user a PPTP server in a RADIUS user profile. In addition, the terminal server features a new command that enables you to route a terminal-server session to a PPTP server. Further, you can specify a host name or an IP address when specifying the route to a PPTP server, and the default value is null. Previously, you could specify only an IP address, and the default value was 0.0.0.0. Finally, a minor change in this release allows a user to delay requesting a PPP session.

Creating tunnels on a per-user basis

In previous releases, when a client dialed into the MAX and wanted to use a PPTP tunnel, the MAX chose a tunnel on the basis of the Route Line *n* parameters. Each T1 PRI line was associated with a different Route Line *n* parameter. Each parameter specified a particular PPTP server at the end of the PPTP tunnel. The MAX simply created a tunnel for each T1 line on which the user connected.

While you can still use the Route Line *n* parameters to create tunnels on the basis of the T1 line, you can now create a tunnel on a per-user basis as well. In a RADIUS user profile, you specify the IP address or host name of a PPTP server. The profile creates a tunnel between the MAX and the PPTP server. When the name and password of an incoming call match the name and password in a RADIUS user profile set up for PPTP, the MAX creates the PPTP tunnel to the PPTP server.

The changes to PPTP functionality affect PPP connections and terminal-server users. This release includes the following new RADIUS attributes:

- Tunnel-Type (64)
- Tunnel-Medium-Type (65)
- Tunnel-Server-Endpoint (66)
- Tunnel-Client-Endpoint (67)
- Tunnel-ID (68)

Routing a terminal-server session to a PPTP server

You can now use the PPTP command in the terminal-server interface to route the session to a PPTP server. This new command gives you two options for selecting the tunnel the MAX

creates. You can specify either the IP address or host name of the PPTP server. Normal PPTP authentication proceeds once the MAX creates the tunnel.

Specifying a host name

For each Route Line *n* parameter in the Ethernet > Mod Config > PPTP Options submenu, you can specify either an IP address or a symbolic hostname.

Other changes

In the past, a problem occurred if the user connected with a modem and didn't use PPP right away. The MAX would wait a short time, and then connect the user to the terminal server, bypassing the PPTP code altogether. The present release solves this problem.

At this time, you cannot use a Connection profile to create a PPTP tunnel on a per-user basis.

Note: New RADIUS attributes for enabling call routing to PPTP provide yet another mechanism for creating PPTP tunnels.

New RADIUS attributes

To enable selection of PPTP tunnels on a per-user basis, this release implements new RADIUS attributes 64–68.

Tunnel-Type (64)

Description: Specifies the type of tunnel the profile can set up.

Usage: At present, you can specify PPTP (1). This setting indicates Point to Point Tunneling Protocol. PPTP enables you to tunnel data between Front End Processors (FEPs) and Windows NT servers, encapsulating PPP frames in enhanced Generic Routing Encapsulation (GRE) packets. Using PPTP, remote users can employ workstations running Windows 95 and Windows NT to dial into a local ISP and connect to their corporate network.

See Also: Tunnel-Medium-Type (65), Tunnel-Server-Endpoint (66), Tunnel-Client-Endpoint (67), Tunnel-ID (68)

Tunnel-Medium-Type (65)

Description: Specifies the type of network on which the MAX builds the tunnel.

Usage: At present, you can specify IP. This setting indicates an IP (Internet Protocol) network.

See Also: Tunnel-Type (64), Tunnel-Server-Endpoint (66), Tunnel-Client-Endpoint (67), Tunnel-ID (68)

Tunnel-Server-Endpoint (66)

Description: Specifies the IP address or hostname of the server at the remote end of the tunnel.

RADIUS features

Tunnel-Client-Endpoint (67)

Usage: Specify an IP address in dotted decimal notation, or a symbolic hostname. The default value is null.

See Also: Tunnel-Type (64), Tunnel-Medium-Type (65), Tunnel-Client-Endpoint (67), Tunnel-ID (68)

Tunnel-Client-Endpoint (67)

Description: Specifies a name or number with which the PPTP server identifies the caller. This attribute provides a way for the PPTP server to keep track of connections.

Usage: Enter the value specified by the PPTP server administrator. The default value is null.

See Also: Tunnel-Type (64), Tunnel-Medium-Type (65), Tunnel-Server-Endpoint (66), Tunnel-ID (68)

Tunnel-ID (68)

Description: Specifies an identifier that designates the tunnel for accounting purposes.

Usage: This attribute is not set in a user profile. The default value is null.

See Also: Tunnel-Type (64), Tunnel-Medium-Type (65), Tunnel-Server-Endpoint (66), Tunnel-Client-Endpoint (67)

RADIUS example

The following user profile sets up a PPTP tunnel over an IP network to the server at IP address 200.168.6.199:

```
mlw Password="pipeline"
      Tunnel-Server-Endpoint="200.168.6.199",
      Tunnel-Type=PPTP,
      Tunnel-Medium-Type=IP
```

New terminal server command

The new PPTP command enables a terminal-server user to use PPP to communicate with a PPTP server. To use the command, enter the following specification at the terminal-server prompt:

pptp *pptp_server*

For the **pptp_server** argument, specify the IP address or symbolic hostname of the PPTP server. When you enter the command, the system displays the following text:

PPTP: Starting session

PPTP Server *pptp_server*

Changes to the Route Line parameter

The following changes apply to the Route Line parameter in the Ethernet > Mod Config > PPTP Options menu:

- You can enter either an IP address in dotted decimal notation or a symbolic hostname. Formerly, you could enter only an IP address.
- The default value is null. Formerly, the default value was 0.0.0.0.

RADIUS support for IPX call and data filters

In previous releases, you could set up data and call filters in RADIUS for IP and generic packets only. Now, you can use RADIUS to create data and call filters for IPX packets as well.

Changes to RADIUS attributes

The Ascend-Data-Filter and Ascend-Call-Filter attributes now support additional arguments for IPX filters.

Here is the syntax for creating an IPX data filter:

```
Ascend-Data-Filter="ipx <dir> <action>
    [srcipxnet <srcipxnet> srcipxnode <srcipxnode>
    [srcipxsoc <cmp> <value> ]]
    [dstipxnet <dstipxnet> dstipxnode <dstipxnode>
    [dstipxsoc <cmp> <value> ]]
```

Here is the syntax for creating an IPX call filter:

```
Ascend-Call-Filter="ipx <dir> <action>
    [srcipxnet <srcipxnet> srcipxnode <srcipxnode>
    [srcipxsoc <cmp> <value> ]]
    [dstipxnet <dstipxnet> dstipxnode <dstipxnode>
    [dstipxsoc <cmp> <value> ]]
```

Note: A filter definition cannot contain newlines. The syntax is shown on multiple lines for documentation purposes only.

Table 26 lists each keyword and argument.

Table 26. IPX filter syntax elements

Syntax element	Description
ipx	Designates an IPX filter.
<dir>	Indicates filter direction. You can specify “in” (to filter packets coming into the MAX) or “out” (to filter packets going out of the MAX).

RADIUS features

RADIUS support for IPX call and data filters

Table 26. IPX filter syntax elements (continued)

Syntax element	Description
<action>	Indicates the action the MAX should take with a packet that matches the filter. You can specify either “forward” or “drop”.
srcipxnet	Designates that a source IPX network number appears after this keyword.
<srcipxnet>	Specifies the source IPX network number—the unique internal network number assigned to the NetWare server. You must specify the network number in hexadecimal format. Specifying 0x or 0X is optional.
srcipxnode	Designates that a source IPX node number appears after this keyword.
<srcipxnode>	Specifies the source IPX node number—the node number of the NetWare server. A valid IPX node number must accompany the IPX network number. You must specify the node number in hexadecimal format. Specifying 0x or 0X is optional. The IPX node number 0xffffffff is allowed and matches all IPX packets with the same node number.
srcipxsoc	Designates that a source IPX socket number specification appears after this keyword.
<cmp>	Indicates how to compare the socket number specified by <value> to the actual socket number in the packet. The <cmp> argument can have the value <, =, >, or !=.
<value>	<p>Specifies the socket number of the NetWare server. Following the srcipxsoc keyword, the <value> argument specifies the source socket number; following the dstipxsoc keyword, the <value> argument specifies the destination socket number.</p> <p>You must specify the socket number in hexadecimal format. Specifying 0x or 0X is optional.</p>
dstipxnet	Designates that a destination IPX network number appears after this keyword.
<dstipxnet>	Specifies the destination IPX network number—the unique internal network number assigned to the NetWare server. You must specify the network number in hexadecimal format. Specifying 0x or 0X is optional.
dstipxnode	Designates that a destination IPX node number appears after this keyword.
<dstipxnode>	<p>Specifies the destination IPX node number—the node number of the NetWare server. A valid IPX node number must accompany the IPX network number.</p> <p>You must specify the node number in hexadecimal format. Specifying 0x or 0X is optional. The IPX node number 0xffffffff is allowed and matches all IPX packets with the same node number.</p>

Table 26. IPX filter syntax elements (continued)

Syntax element	Description
dstipxsoc	Designates that a source IPX socket number specification appears after this keyword.

IPX filter examples

Example 1

The IPX filter specified in the following RADIUS user profile drops all outbound IPX packets with a destination IPX network number of 0x00003823, regardless of the node or socket number. The generic filter that appears after the IPX filter forwards all other packets.

```
st1 Password="st1"
    Ascend-Idle-Limit=300,
    Ascend-Route-IPX=Route-IPX-Yes,
    Ascend-Route-IP =Route-IP-Yes,
    Ascend-IPX-Peer-Mode=IPX-Peer-Router,
    Ascend-Data-Filter="ipx out drop dstipxnet 0x00003823 dstipxnode
0xffffffffffff",
    Ascend-Data-Filter="generic out forward 0 0 0"
```

You should specify a default filter for packets that do not match the filter criteria. In this example, if the specification Ascend-Data-Filter ="generic out forward 0 0 0" did not appear in the profile, the MAX would drop all other IPX, IP, and generic packets.

Example 2

The IPX filter specified in the following RADIUS user profile drops all outbound IPX packets with a source network number of 0x00000005, a source node number of 00abcde12345, and a source socket number of 0x4002. The generic filter that appears after the IPX filter forwards all other packets.

```
st1 Password = "st1"
    Ascend-Idle-Limit=300,
    Ascend-Route-IPX=Route-IPX-Yes,
    Ascend-Route-IP =Route-IP-Yes,
    Ascend-IPX-Peer-Mode=IPX-Peer-Router,
    Ascend-Data-Filter="ipx in drop srcipxnet 0x00000005 srcipxnode
0x00abcde12345
00a024cd5807 srcipxsock=0x4002",
    Ascend-Data-Filter="generic out forward 0 0 0"
```

Note: A filter definition cannot contain newlines. The syntax is shown on multiple lines for documentation only.

Modem ID added to RADIUS Accounting Stop Records

You can now get modem slot and port information through RADIUS Accounting Stop records.

Background

Previously, RADIUS Accounting records included detailed statistics about time connected, bandwidth used, and the user logged in, but no specific information about the MAX modem in use for the session.

RADIUS accounting Stop records now provide information about the specific modem that each client uses. You can use the information to monitor the status of each call into a specific modem, and to generate statistics based on call disconnects, progress codes, call duration, and the modem used for the call.

New RADIUS attributes

Two new RADIUS attributes support modem identification in Accounting Stop records. Ascend-Modem-PortNo (Attribute 120) specifies the modem used for the call, and Ascend-Modem-SlotNo (Attribute 121) specifies the slot that holds the modem.

Ascend-Modem-PortNo (Attribute 120)

Description: Specifies, for inclusion in an accounting Stop record, the modem used for the call.

Usage: The MAX sends Ascend-Modem-PortNo as part of an accounting Stop record. The attribute does not appear in a user profile.

Dependencies: Because the MAX designates a modem by slot card and port, you must consider the value of Ascend-Modem-SlotNo.

See Also: Ascend-Modem-SlotNo

Ascend-Modem-SlotNo (Attribute 121)

Description: Specifies, for inclusion in an accounting Stop record, the slot containing the modem used for the call.

Usage: The MAX sends Ascend-Modem-SlotNo as part of an accounting Stop record. The attribute does not appear in a user profile.

Dependencies: Because the MAX designates a modem by slot card and port, you must consider the value of Ascend-Modem-PortNo

See Also: Ascend-Modem-PortNo

RADIUS now follows NI-2 PRI spec for outgoing calls

A new RADIUS attribute defines the call numbering plan and enables the MAX to place a 7-digit local call to an NI-2 switch using standard ISDN numbering. Prior to this release, the MAX had to place this type of call as a 10-digit call (long distance call).

How it works

A new RADIUS attribute, Ascend-Numbering-Plan-ID, defines the numbering plan ID for outbound PRI calls. Ascend-Numbering-Plan-ID works with the RADIUS attribute Ascend-PRI-Number-Type to describe the number the outgoing call dials. The following combinations are permitted by the NI-2 standard:

Ascend-PRI-Numbering-Type	Ascend-Numbering-Plan-ID
International-Number (1)	ISDN-Numbering-Plan(1)
National-Number(2)	ISDN-Numbering-Plan(1)
Local-Number(4)	ISDN-Numbering-Plan(1)
Local-Number(4)	Private-Numbering-Plan(9)
Unknown-Number(0)	Unknown-Numbering-Plan(0)

To set up a PRI call to an NI-2 switch, you could use the following combination in the RADIUS user profile:

- Ascend-PRI-Numbering-Type=Unknown-Number
- Ascend-Numbering-Plan-ID=Unknown-Numbering-Plan

With this combination the switch routes the call based only upon the actual digits dialed. If you dial only seven digits, the switch assumes the area code for the number dialed is the same as the number dialed from. You would dial 202-232-2232 to reach the same number if you are calling from a number within the local toll-free calling area, and 1-202-232-2232 when calling from anywhere else.

If you use a combination for a local call other than Ascend-PRI-Numbering-Type=Unknown-Number and Ascend-Numbering-Plan-ID=Unknown-Numbering-Plan, the switch depends upon you to determine how the call is to be routed. For example, a PBX can use this feature to do least-cost routing to multiple long distance providers.

Ascend-Numbering-Plan-ID

Description: Ascend-Numbering-Plan-ID (105) enables you to select the PRI for outbound PRI calls.

Usage: To set up a PRI call to an NI-2 switch, try the following user profile first:

- Ascend-PRI-Numbering-Type=Unknown-Number
- Ascend-Numbering-Plan-ID=Unknown-Numbering-Plan

SNMP features

SNMP can enable/disable/quiesce T1/PRI links

This combination of settings should allow you to dial any number, long-distance, local, or local area toll-free, since the NI-2 switch routes the call based upon the numbers actually dialed. Any other combination of settings may require additional network configuration.

Note: Some NI-2 switches may prefer another combination of number type and numbering plan. See the Dependencies paragraph, below.

The possible values for Ascend-Numbering-Plan-ID are:

- **Unknown-Numbering-Plan (0)**
Allows you to dial any number (7-digit, 9-digit, or 10-digit) for an outgoing PRI call. When Ascend-Numbering-Plan-ID=Unknown the Ascend-PRI-Number-Type for the same user profile should be Unknown.
- **ISDN-Numbering-Plan (1)**
Specifies that the MAX use the NI-2 standard numbering plan, E.164. You cannot use this value with Ascend-PRI-Numbering-Type=Unknown-Number if the outgoing call is to an NI-2 switch, since the NI-2 specification does not permit this combination. ISDN-Numbering-Plan is the default.
- **Private-Numbering-Plan (9)**
Specifies a numbering plan unique to the number being called. For example, a company with a 5-digit internal dialing plan might use a private numbering plan. In all cases, the number dialed should agree with the number type. For example, it would not work to set the number type to International and then dial a 7-digit number.

Dependencies: The outgoing call must to an NI-2 switch.must conform to NI-2 requirements with respect to the settings. The permitted setting pairs are:

Ascend-PRI-Numbering-Type	Ascend-Numbering-Plan-ID
International-Number (1)	ISDN-Numbering-Plan(1)
National-Number(2)	ISDN-Numbering-Plan(1)
Local-Number(4)	ISDN-Numbering-Plan(1)
Local-Number(4)	Private-Numbering-Plan(9)
Unknown-Number(0)	Unknown-Numbering-Plan(0)

See Also: Ascend-PRI-Numbering-Type

SNMP features

SNMP can enable/disable/quiesce T1/PRI links

SNMP management applications can now change the configuration state of T1/PRI links.

Changes to the Ascend MIB

A new SNMP variable, wanLineUsage (1.3.6.1.4.1.529.4.21.1.8), enables an SNMP management application to set the configuration state of a T1 line. Although the MIB object includes other status options, only the following are currently supported for SNMP set commands:

- lu-disabled(3)
- lu-trunk(5)
- lu-quiesced(6)

```
wanLineUsage OBJECT-TYPE
    SYNTAX  INTEGER {
        lu-unknown(1),
        lu-unavailable(2),
        lu-disabled(3),
        lu-enabled(4),
        lu-trunk(5),
        lu-quiesced(6),
        lu-drop-and-insert(7),
        lu-t-online-user(8),
        lu-t-online-zgr(9)
    }
    ACCESS  read-write
    STATUS  mandatory
    DESCRIPTION
        The usage of the line. SETting values are supported only
for
        1/E1 links. Only lu-disabled(3), lu-trunk(5), and
        lu-quiesced(6) for T1/PRI are currently supported for
SETs. "
 ::= { wanLineEntry 8 }
```

Note: SNMP can obtain the status of a T1/PRI line using wanLineUsage.

Enable and Disable individual modem using SNMP

A new MIB object makes it possible for you to enable or disable an individual modem in a T1 MAX using SNMP.

Changes to the Ascend MIB

You can enable or disable individual modems on MAX units with the T1 interface platform feature using the new MIB object slotMdmItemConfig (1.3.6.1.4.1.529.2.5.1.6). The modem slot card must be enabled in order to enable or disable the modem.

The following was added to the slot modem table for slotMdmEntry.

```
slotMdmItemConfig OBJECT-TYPE

    SYNTAX  INTEGER {
```

SNMP features

SNMP write security disabled by default

```
        other(1),
        enable(2),
        disable(3),
        disableAndChannel(4)
    }
ACCESS      read-write
STATUS      mandatory
DESCRIPTION "The modem configuration state. SETs
            are allowed only if the corresponding
            modem slot card is enabled."
::= { slotMdmEntry 6 }
```

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 an Ascend unit's system 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 > 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: This parameter 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.

SNMP can help find which device a call is logged into

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-1

Hunt-2#

Hunt-3

Description: These parameters indicate the hunt group numbers 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

Location: Net T1 Line Profile > Line Config

Ascend MIB changes

New variables have been added to the wanLineTable in the Ascend MIB, wan.mib, to support this feature:

```
WanLineEntry ::=
    SEQUENCE {
```

SNMP features

Hunt-1 # Hunt-2# Hunt-3 #

```
        wanLineIfIndex
            INTEGER,
        wanLineName
            DisplayString,
        wanLineType
            OBJECT IDENTIFIER,
        wanLineChannels
            INTEGER,
        wanLineState
            INTEGER,
        wanLineStateString
            DisplayString,
        wanLineActiveChannels
            INTEGER,
        wanLineUsage
            INTEGER,
        wanLineHuntGrpPhoneNumber1
            DisplayString,
        wanLineHuntGrpPhoneNumber2
            DisplayString,
        wanLineHuntGrpPhoneNumber3
            DisplayString
    }

wanLineHuntGrpPhoneNumber1 OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-write
    STATUS  mandatory
    DESCRIPTION
        "The hunt group phone number associated with
        the line. This entry is manually entered in the line
        configurations options."
    ::= { wanLineEntry 9 }

wanLineHuntGrpPhoneNumber2 OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-write
    STATUS  mandatory
    DESCRIPTION
        "The hunt group phone number associated with
        the line. This entry is manually entered in the line
        configurations options."
    ::= { wanLineEntry 10 }

wanLineHuntGrpPhoneNumber3 OBJECT-TYPE
    SYNTAX  DisplayString
    ACCESS  read-write
    STATUS  mandatory
    DESCRIPTION
        "The hunt group phone number associated with
        the line. This entry is manually entered in the line
        configurations options."
    ::= { wanLineEntry 11 }
```

SNMP "get" now retrieves MPP session statistics

A new table in the Ascend MIB, mppActiveStatsTable, enables you to use an SNMP get request to obtain MPP the session statistics that appear in the Dyn Stat status display.

Overview

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 Window Parameter	Description
mppStatsRemoteName	<i>(Connection Profile name)</i>	Connection Profile name set up in the MAX 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)
mppStatsStartingTimeSta mp	<i>(time)</i>	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.

SNMP features

SNMP "get" now retrieves MPP session statistics

Value in mppStatsTable	Dyn Stats Window Parameter	Description
mppStatsBandwidth	(data rate)	Third line of Dyn Stats window. Shows the current data rate
mppStatsTotalChannels	<i>n</i> channels	The number of channels the data rate in mppStatsBandwidth represents.
mppStatsCLU	CLU <i>n</i> %	Current line utilization
mppStatsALU	ALU <i>n</i> %	Average line utilization.

Changes to the Ascend MIB

A new table, the mppActiveStatsTable, has been added to the SessionStatusGroup in the Ascend MIB to support this feature.

- **mppActiveStatsTable** OBJECT-TYPE sessionStatusGroup 4
SYNTAX SEQUENCE OF MppActiveStatsEntry
ACCESS not-accessible
STATUS mandatory
DESCRIPTION "A list of active MPP session statistics with invalid entries screened out and indexed by mppStatsMpID."
- mppActiveStatsEntry OBJECT-TYPE mppActiveStatsTable 1
SYNTAX MppActiveStatsEntry
ACCESS not-accessible
STATUS mandatory
DESCRIPTION "An entry containing object variables to describe an active MPP session. The variables are those seen in the Dyn Stat area of the LCD display."
- **mppStatsMpID** OBJECT-TYPE mppActiveStatsEntry 1
SYNTAX INTEGER
ACCESS read-only
STATUS mandatory
DESCRIPTION "The MpID number for this active MPP session entry."
- mppStatsRemoteName OBJECT-TYPE mppActiveStatsEntry 2
SYNTAX DisplayString
ACCESS read-only
STATUS mandatory
DESCRIPTION "The name of the remote user."
- mppStatsQuality OBJECT-TYPE mppActiveStatsEntry 3
SYNTAX INTEGER {
good(1),
fair(2),
marginal(3),
poor(4),

na(5)

ACCESS read-only

STATUS mandatory

DESCRIPTION Line quality. N/A: No MPP sessions currently active,

Good: <%1 CRC errors,

Fair: <%5 CRC errors,

Marginal: <%10 CRC errors,

Poor:%10 or > CRC errors"

- **mppStatsBandwidth** OBJECT-TYPE mppActiveStatsEntry 4
SYNTAX INTEGER
ACCESS read-only
STATUS mandatory
DESCRIPTION "Total bit rate (Kbps) for the MPP session."
- **mppStatsTotalChannels** OBJECT-TYPE mppActiveStatsEntry 5
SYNTAX INTEGER
ACCESS read-only
STATUS mandatory
- **DESCRIPTION** "The total number of channels associated with this MPP session."
- **mppStatsCLU** OBJECT-TYPE mppActiveStatsEntry 6
SYNTAX INTEGER
ACCESS read-only
STATUS mandatory
DESCRIPTION "Current percentage of line utilization for transmitted packets during this MPP session."
- **mppStatsALU** OBJECT-TYPE mppActiveStatsEntry 7
SYNTAX INTEGER
ACCESS read-only
STATUS mandatory
DESCRIPTION "Average percentage of line utilization for transmitted packets during this MPP session."
- **mppStatsStartingTimeStamp** OBJECT-TYPE mppActiveStatsEntry 8
SYNTAX INTEGER
ACCESS read-only
STATUS mandatory

DESCRIPTION "The starting time for this MPP session in seconds since startup."

SNMP can monitor WAN lines and channels

You can now use SNMP to monitor WAN lines without logging into the MAX.

Overview

You can use SNMP to obtain WAN information without logging into the MAX. Each Ascend WAN type is assigned an object ID. These identifiers are the root of the MIB subtree containing WAN-specific information. These subtrees, when appropriate, are described in separate files.

Changes to the Ascend MIB

- wanUseTrunkGroups OBJECT-TYPE wanInfo 20
SYNTAX INTEGER {
do-not-use(1),
use(2)
ACCESS read-only
STATUS mandatory
DESCRIPTION "System wide parameter dictating the use of trunk groups."
- wanLineTable OBJECT-TYPE wanInfo 21
SYNTAX SEQUENCE OF WanLineEntry
ACCESS not-accessible
STATUS mandatory
DESCRIPTION "The wan line table."
- wanLineEntry OBJECT-TYPE wanLineTable 1
SYNTAX WanLineEntry
ACCESS not-accessible
STATUS mandatory
DESCRIPTION "An entry in the wan line table."
- wanLineIfIndex OBJECT-TYPE wanLineEntry 1
SYNTAX INTEGER
ACCESS read-only
STATUS mandatory
DESCRIPTION "This value for this object is equal to the value of ifIndex from the Interfaces table of MIB II (RFC 1213)."
- wanLineName OBJECT-TYPE wanLineEntry 2
SYNTAX DisplayString
ACCESS read-only
STATUS mandatory
DESCRIPTION "A textual name of the wanLine as assigned through the menu sytem."
- wanLineType OBJECT-TYPE wanLineEntry 3
SYNTAX OBJECT IDENTIFIER
ACCESS read-only
STATUS mandatory

DESCRIPTION "One of 'wanTypes'."

- wanLineChannels OBJECT-TYPE wanLineEntry 4

SYNTAX INTEGER

ACCESS read-only

STATUS mandatory

DESCRIPTION "The number of ds0 channels supported."

- wanLineState OBJECT-TYPE wanLineEntry 5

SYNTAX INTEGER {

ls-unknown(1),

ls-does-not-exist(2),

ls-disabled(3),

ls-no-physical(4),

ls-no-logical(5),

ls-point-to-point(6),

ls-multipoint-1(7),

ls-multipoint-2(8),

ls-loss-of-sync(9),

ls-yellow-alarm(10),

ls-ais-receive(11),

ls-no-d-channel(12),

ls-active(13),

ls-maintenance(14)

ACCESS read-only

STATUS mandatory

DESCRIPTION "The state of the line."

- wanLineStateString OBJECT-TYPE wanLineEntry 6

SYNTAX DisplayString

ACCESS read-only

STATUS mandatory

DESCRIPTION "A textual representation of the wanLineState as displayed by the menu sytem."

- wanLineActiveChannels OBJECT-TYPE wanLineEntry 7

SYNTAX INTEGER

ACCESS read-only

STATUS mandatory

DESCRIPTION "The number of active ds0 channels of the line."

- wanLineChannelTable OBJECT-TYPE wanInfo 22

SYNTAX SEQUENCE OF WanLineChannelEntry

ACCESS not-accessible

STATUS mandatory

SNMP features

SNMP can monitor WAN lines and channels

DESCRIPTION "The wan line table."

- wanLineChannelEntry OBJECT-TYPE wanLineChannelTable 1

SYNTAX WanLineChannelEntry

ACCESS not-accessible

STATUS mandatory

DESCRIPTION "An entry in the wan line table."

- wanLineChannelIfIndex OBJECT-TYPE wanLineChannelEntry 1

SYNTAX INTEGER

ACCESS read-only

STATUS mandatory

DESCRIPTION "This value for this object is equal to the value of ifIndex from the Interfaces table of MIB II (RFC 1213)."

- wanLineChannelIndex OBJECT-TYPE wanLineChannelEntry 2

SYNTAX INTEGER

ACCESS read-only

STATUS mandatory

DESCRIPTION "The ds0 channel number with the line."

- wanLineChannelState OBJECT-TYPE wanLineChannelEntry 3

SYNTAX INTEGER {

bs-unknown(1),

bs-unavailable(2),

bs-unused(3),

bs-out-of-service(4),

bs-nailed-up(5),

bs-held(6),

bs-idle(7),

bs-clear-pending(8),

bs-dialing(9),

bs-ringing(10),

bs-connected(11),

bs-signaling(12),

bs-cut-through(13),

bs-current-d(14),

bs-backup-d(15),

bs-maintenance(16),

bs-spc-up(17)

ACCESS read-only

STATUS mandatory

DESCRIPTION "The state of the ds0 channel."

- wanLineChannelStateString OBJECT-TYPE wanLineChannelEntry 4
SYNTAX DisplayString
ACCESS read-only
STATUS mandatory
DESCRIPTION "A textual representation of the wanLineChannelState as displayed by the menu sytem."
- wanLineChannelErrorCount OBJECT-TYPE wanLineChannelEntry 5
SYNTAX INTEGER
ACCESS read-only
STATUS mandatory
DESCRIPTION "The error count encountered on the channel."
- wanLineChannelUsage OBJECT-TYPE wanLineChannelEntry 6
SYNTAX INTEGER {
ds0-unused-channel(1),
ds0-switched-channel(2),
ds0-cut-through(3),
ds0-clear-64(4),
ds0-pri-d-channel(5),
ds0-nfas-prime-d(6),
ds0-nfas-sec-d(7),
ds0-cas-channel(8),
ds0-spc-channel(9)
ACCESS read-only
STATUS mandatory
DESCRIPTION "The usage for this ds0 channel."
- wanLineChannelTrunkGroup OBJECT-TYPE wanLineChannelEntry 7
SYNTAX INTEGER
ACCESS read-only
STATUS mandatory
DESCRIPTION "The trunk group assigned to this channel."
- wanLineChannelPhoneNumber OBJECT-TYPE wanLineChannelEntry 8
SYNTAX DisplayString
ACCESS read-only
STATUS mandatory
DESCRIPTION "The phone number of this channel. This is the number sent to the far end in an inverse multiplexed call when instructing the far end to add more bandwidth. The number should contain the minimum number of digits to identify the channel. If the channel is part of a hunt group, the phone number should be blank."
- wanLineChannelSlot OBJECT-TYPE wanLineChannelEntry 9
SYNTAX INTEGER

SNMP features

SNMP can obtain active call status

ACCESS read-only

STATUS mandatory

DESCRIPTION "A slot number for routing incoming calls associated with the channel. A slot-port number zero means calls arriving on this channel can be routed to any port."

- wanLineChannelPort OBJECT-TYPE wanLineChannelEntry 10

SYNTAX INTEGER

ACCESS read-only

STATUS mandatory

DESCRIPTION "A port number for routing incoming calls associated with the channel. A slot-port number zero means calls arriving on this channel can be routed to any port."

- wanLineChannelNailedState OBJECT-TYPE wanLineChannelEntry 11

SYNTAX INTEGER {
not-applicable(1),
nailed-held(2),
nailed-active(3)

ACCESS read-only

STATUS mandatory

DESCRIPTION "The nailed group associated with the channel."

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..'7ffffff'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 "Differentiates between outgoing and incoming calls."

SNMP system reset

You can now reset an Ascend unit using SNMP.

Overview

The new SNMP object `sysReset` enables you to reset an Ascend unit from an SNMP manager. A one-minute timeout after the reset command is issued permits the Ascend unit to confirm the set request before the unit is reset. While the unit is in the process of resetting, it will not respond to other SNMP requests.

Note: You cannot modify the length of the timeout.

Determining whether the Ascend unit has reset

Information held in the Ascend Events Group is erased and its values are initialized when the Ascend unit is reset by software or by toggling the power off and on. `sysAbsoluteStartupTime` is the time in seconds since January 1, 1990. You can retrieve `sysAbsoluteStartupTime` and compare this value against the previous polls value to determine whether the box has actually been reset by `sysReset`.

Traps generated

The reset process generates the standard traps (see `coldStart` and `warmStart` traps in the *SNMP Supplement*).

Changes to Ascend MIB

A new object has been added to the `systemStatusGroup` in the Ascend MIB:

<code>sysReset</code>	OBJECT-TYPE
SYNTAX	INTEGER { no-op(1), reset(2) }
ACCESS	read-write
STATUS	mandatory
DESCRIPTION	"The reset takes effect after 1 minute."
::= {	systemStatusGroup 8 }

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 `ssnStatusCallReferenceNum` from 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

SNMP features

SNMP can detect concurrent sessions

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."

ASCEND MIB now includes modem status objects

You can now use SNMP to monitor modem .status for the modems in an 8-modem or 12-modem card. The parameters you can monitor are the same as those displayed in the Modem Status window.

Ascend MIB additions

The Main Status Menu now contains a V.34 Modem status entry for each individual modem on an 8-modem card and a 12-modem card. When you select the V.34 Modem entry for a card, the Modem Status menu displays. This feature contains additions to the Ascend MIB that enable you to monitor modem status using SNMP.

The following additions were made to the `slotMdmEntry` group in the Ascend MIB.

`slotMdmEntry`

```
slotMdmItemStatusOBJECT-TYPE
SYNTAX INTEGER {
    modemStateNonExist( 1 ),
    modemStateFailPost( 2 ),
    modemStateIdle( 3 ),
    modemStateAwaitingRlsd( 4 ),
    modemStateAwaitingCodes( 5 ),
    modemStateOnline( 6 ),
```

```

modemStateInit( 7 ),
modemStateInitOpenQueued( 8 ),
modemStateInitOpenQueuedVC( 9 ),
modemStateInitDialStr2( 10 ),
modemStateInitDialStr3( 11 ),
modemStateVirtualConnect( 12 ),
modemStateDisabled( 13 ),
modemStateDisabledChan( 14 )
}

ACCESS read-only
STATUS mandatory
DESCRIPTION "The status of the modem."
::= { slotMdmEntry 4 }

slotMdmItemStatusChar OBJECT-TYPE
SYNTAX DisplayString
ACCESS read-only
STATUS mandatory
DESCRIPTION "The status of the modem as displayed in the menu system."
::= { slotMdmEntry 5 }

```

SNMP RFC 1398 Ethernet-like MIB (dot3) support added

Support has been added to the the MIB for monitoring statistics for Ethernet-like objects, conforming to RFC 1398. This support includes counters for specific types of frames and errors on a per-interface basis.

RFC 1398 is the mib for the Ethernet port in Ascend units.

Set system clock through SNMP

A new object in the Ascend MIB enables you to use SNMP to set the system clock.

The Ascend MIB now includes the following object:

```

sysAbsoluteCurrentTime OBJECT-TYPE
SYNTAX      INTEGER (1..'7fffffff'h)
ACCESS      read-write
STATUS      mandatory
DESCRIPTION  "The current system time in seconds since
              January 1, 1990. Changing this value may
              result in a change of sysAbsoluteStartupTime
              and not of sysSecsSinceStartup. The following
              relationship holds:

              sysAbsoluteCurrentTime -
              sysAbsoluteStartupTime = sysSecsSinceStartup."
::= { systemStatusGroup 7 }

```


Terminating user sessions using SNMP

You can now terminate a user session using SNMP. You can now set the `ssnStatusValidFlag` object in the Session Status group of the Ascend MIB. To terminate a user session, set `ssnStatusValidFlag` to invalid (1).

The new description for `ssnStatusValidFlag` is presented below.

```
ssnStatusValidFlag OBJECT-TYPE
    SYNTAX INTEGER {
        invalid(1),
        valid(2)
    }
    ACCESS read-write
    STATUS mandatory
    DESCRIPTION "Value indicates whether this entry contains
valid
        information or not. Setting a valid(2) session to
invalid(1)
        causes an effective session termination."
 ::= { sessionStatusEntry 2 }
```

Ascend MIB change supports counters for total and current calls

A change has been made to the Ascend MIB to permit service management (including capacity planning) based upon the number of digital/analog/frame relay calls for each MAX.

How the MIB changes work

This feature changes the Ascend MIB to support counters for incoming and outgoing calls. These counters track the current and total number of digital, analog, and frame relay calls for each MAX. These counters are reset and initialized to zero when the MAX is reset or booted up.

Changes to Ascend MIB

The changes below have been made to the `callStatus` group to support the call counters.

```
callCurrentAnalogOutgoing OBJECT-TYPE
    SYNTAX INTEGER
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION "The number of current analog outgoing calls is
returned."
 ::= { callStatusGroup 4 }

callCurrentAnalogIncoming OBJECT-TYPE
```

SNMP features

Ascend MIB change supports counters for total and current calls

```
SYNTAXINTEGER
ACCESSread-only
STATUSmandatory
DESCRIPTION"The number of current analog incoming calls is
returned."
::= { callStatusGroup 5 }

callCurrentDigitalOutgoing OBJECT-TYPE
SYNTAXINTEGER
ACCESSread-only
STATUSmandatory
DESCRIPTION"The number of current digital outgoing calls is
returned."
::= { callStatusGroup 6 }

callCurrentDigitalIncoming OBJECT-TYPE
SYNTAXINTEGER
ACCESSread-only
STATUSmandatory
DESCRIPTION"The number of current digital incoming calls is
returned."
::= { callStatusGroup 7 }

callCurrentFROutgoing OBJECT-TYPE
SYNTAXINTEGER
ACCESSread-only
STATUSmandatory
DESCRIPTION"The number of current frame relay outgoing calls
is returned."
::= { callStatusGroup 8 }

callCurrentFRIncoming OBJECT-TYPE
SYNTAXINTEGER
ACCESSread-only
STATUSmandatory
DESCRIPTION"The number of current frame relay incoming calls
is returned."
::= { callStatusGroup 9 }

callTotalAnalogOutgoing OBJECT-TYPE
SYNTAXINTEGER
ACCESSread-write
STATUSmandatory
DESCRIPTION"The total number of analog outgoing calls since
system bootup or last clear of the variable is returned."
::= { callStatusGroup 10 }

callTotalAnalogIncoming OBJECT-TYPE
SYNTAXINTEGER
ACCESSread-write
STATUSmandatory
DESCRIPTION"The total number of analog incoming calls since
```

```
system bootup or last clear of the variable is returned."
 ::= { callStatusGroup 11 }

    callTotalDigitalOutgoing OBJECT-TYPE
    SYNTAXINTEGER
    ACCESSread-write
    STATUSmandatory
    DESCRIPTION"The total number of digital outgoing calls since
system bootup or last clear of the variable is returned."
 ::= { callStatusGroup 12 }

    callTotalDigitalIncoming OBJECT-TYPE
    SYNTAXINTEGER
    ACCESSread-write
    STATUSmandatory
    DESCRIPTION"The total number of digital incoming calls since
system bootup or last clear of the variable is returned."
 ::= { callStatusGroup 13 }

    callTotalFROutgoing OBJECT-TYPE
    SYNTAXINTEGER
    ACCESSread-write
    STATUSmandatory
    DESCRIPTION"The total number of frame relay outgoing calls
since system bootup or last clear of the variable is returned."
 ::= { callStatusGroup 14 }

    callTotalFRIncoming OBJECT-TYPE
    SYNTAXINTEGER
    ACCESSread-write
    STATUSmandatory
    DESCRIPTION"The total number of frame relay incoming calls
since system bootup or last clear of the variable is returned."
 ::= { callStatusGroup 15 }
```

Firewall Control Protocol managed by SNMP

SAM firewalls embedded in Ascend products can now be managed through SNMP.

How the Firewall Control Protocol works

The SNMP objects in the sysFcpGroup of the Ascend Enterprise MIB provide a means of enabling and disabling, creating and changing SAM firewalls in the Ascend MIB.

Ascend MIB definitions for the Firewall Control Protocol

```
sysFcpGroup      OBJECT IDENTIFIER ::= { systemStatusGroup 11 }

sysFcpRuleName    OBJECT-TYPE
```

SNMP features

Firewall Control Protocol managed by SNMP

```
SYNTAX          DisplayString
ACCESS          read-write
STATUS          mandatory
DESCRIPTION     "The name of the firewall rule to be
                  enabled or disabled. This name corresponds with a
                  name established when the firewall was created (as
                  by the Secure Access Manager)."
```

```
::= { sysFcpGroup 1 }
```



```
sysFcpExecute   OBJECT-TYPE
SYNTAX          INTEGER {
                  no-op( 1 ),
                  enb-rule( 2 ),
                  dis-rule( 3 )
                }
ACCESS          read-write
STATUS          mandatory
DESCRIPTION     "Cause a firewall given by the above
                  parameters to be affected as requested.
                  add-rule causes a dynamic rule to be created;
                  del-rule causes a dynamic rule to cease operating."
```

```
::= { sysFcpGroup 2 }
```



```
sysFcpTimeOut   OBJECT-TYPE
SYNTAX          INTEGER
ACCESS          read-write
STATUS          mandatory
DESCRIPTION     "Time, expressed in seconds, during which
                  this firewall rule will effect the firewall. After
                  the expiration time, the rule will cease being
                  effective exactly as if a del-rule (see
                  sysFcpExecute above) had been executed on it.
                  Default is 3600 seconds."
```

```
::= { sysFcpGroup 3 }
```



```
sysFcpExtAddr   OBJECT-TYPE
SYNTAX          IpAddress
ACCESS          read-write
STATUS          mandatory
DESCRIPTION     "Address of entity outside firewall. This
                  value defaults to 0.0.0.0, equivalent to
                  a don't care. May be used when selecting
                  firewall to be updated (see sysFcpRoutAddr)."
```

```
::= { sysFcpGroup 4 }
```



```
sysFcpExtAddrMask OBJECT-TYPE
SYNTAX          IpAddress
ACCESS          read-write
STATUS          mandatory
DESCRIPTION     "Netmask of entity outside firewall. This
                  value defaults to 255.255.255.255, equivalent
                  to a host address if sysFcpExtAddr is non-zero."
```

```
::= { sysFcpGroup 5 }

sysFcpExtPort      OBJECT-TYPE
SYNTAX             INTEGER
ACCESS             read-write
STATUS             mandatory
DESCRIPTION        "For external entity, specifies a port number.
                    Defaults to 0, equivalent to don't care."
::= { sysFcpGroup 6 }

sysFcpExtPortMax   OBJECT-TYPE
SYNTAX             INTEGER
ACCESS             read-write
STATUS             mandatory
DESCRIPTION        "For external entity, specifies the maximum
                    port number of a range of ports. Defaults to
                    0, equivalent to specifying a single port number
                    if sysfcpExtPort is nonzero."
::= { sysFcpGroup 7 }

sysFcpIntAddr      OBJECT-TYPE
SYNTAX             IpAddress
ACCESS             read-write
STATUS             mandatory
DESCRIPTION        "Address of entity inside firewall. This
                    value defaults to 0.0.0.0, equivalent to
                    a don't care."
::= { sysFcpGroup 8 }

sysFcpIntAddrMask  OBJECT-TYPE
SYNTAX             IpAddress
ACCESS             read-write
STATUS             mandatory
DESCRIPTION        "Netmask of entity inside firewall. This
                    value defaults to 255.255.255.255, equivalent
                    to a host address if sysFcpIntAddr is non-zero."
::= { sysFcpGroup 9 }

sysFcpIntPort      OBJECT-TYPE
SYNTAX             INTEGER
ACCESS             read-write
STATUS             mandatory
DESCRIPTION        "For Internal entity, specifies a port
                    number. Defaults to 0, equivalent to don't care."
::= { sysFcpGroup 10 }

sysFcpIntPortMax   OBJECT-TYPE
SYNTAX             INTEGER
ACCESS             read-write
STATUS             mandatory
DESCRIPTION        "For Internal entity, specifies the maximum
                    port number of a range of ports. Defaults to
```

SNMP features

SNMP request authentication added

```
0, equivalent to specifying a single port number
if sysFcpIntPort is nonzero."
::= { sysFcpGroup 11 }

sysFcpRoutAddr      OBJECT-TYPE
SYNTAX              IPAddress
ACCESS              read-write
STATUS              mandatory
DESCRIPTION         "This address may be supplied by the
                    management station to choose a firewall for
                    alteration. The default for this address is
                    0.0.0.0, which would cause the router to use
                    sysFcpExtAddr to choose its firewall instead."
::= { sysFcpGroup 12 }

sysFcpAddrOpts      OBJECT-TYPE
SYNTAX              INTEGER
ACCESS              read-write
STATUS              mandatory
DESCRIPTION         "Firewall requests may require additional
                    bit-encoded options to determine the firewall's
                    new behavior. This options variable is a mechanism
                    to allow those options to be defined at a later
                    date."
::= { sysFcpGroup 13 }
```

SNMP request authentication added

This feature adds proprietary SNMP request authentication, including replay protection, to the existing SNMP v1 implemented in Ascend units, to the Ascend MAX and Pipeline products. This implementation of SNMP request authentication is compatible with standard SNMPv1 practices, and affects the Ascend unit's interpretation of SNMP messages that use it. Previously, Ascend units did not provide authentication of SNMP requests.

Authenticating requests using SNMP

You can use SNMP for security-related operations, such as altering the operational state of the Ascend unit (rebooting, loading configurations, etc.) or firewall configurations. This feature adds an authentication option to existing SNMP v1 that verifies 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.

This feature is an addition to SNMP v1 already implemented on Ascend MAX and Pipeline units. You can still use SNMP without authentication by using the previous version of the SNMP R/W Comm parameter in Ethernet > Mod Config > SNMP Options.

Authentication Elements

This feature uses 4 elements to authenticate SNMP packets:

- secret authentication key
- data to be authenticated
- time-dependent state variables (for replay 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 Ascend unit 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

Prior to this software release, existing community names on Ascend units were simply ASCII strings with no internal structure. The original SNMPv1 definition of the community string is a string of octets that is compared to a similar string in the receiving SNMP entity. If the string in the packet received exactly matches a community string in the receiving entity, then the packet is considered "authentic".

The defaults for SNMP v1 (without authentication) are:

Ethernet > Mod Config > SNMP Options > Read Comm=public

Ethernet > Mod Config > SNMP Options > R/W Comm=write

You use a new version of the Read/Write community string if you wish to use SNMP authentication, with the format:

Ethernet > Mod_config > SNMP Options > R/W Comm=write|secretkey

This causes the Ascend unit to require SNMP SET REQUEST packets to be authenticated, using "secretkey" as the shared (but not transmitted) secret.

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**.

Configuring the SNMP manager to use SNMP authentication

To communicate with an Ascend unit that has been configured to use authenticated SNMP, an SNMP management station must construct an SNMP packet using the new format for the Read/Write community string, including the secret key:

name/secretkey

If the Ascend unit has been configured to use authenticated SNMP, it will not accept packets from an SNMP management station using the string format without the pipe/vertical bar.

Initiating RADIUS updates using SNMP

When the Ascend unit powers up, it can retrieve a potentially large quantity of configuration information from the RADIUS server. Some of the data on the RADIUS server can change during operation. Using the Upd Rem Cfg command from the Sys Diag menu, you can instruct the Ascend unit to retrieve a fresh configuration. However, using this command is not convenient if you manage the Ascend unit with an SNMP manager. In this release, you can initiate a RADIUS configuration update using the SNMP Set command, and use SNMP to poll the status of the update.

A new SNMP variable, sysConfigRadiusCmd, enables an SNMP manager to initiate a RADIUS configuration retrieval of routes, IP pools, connection information, and terminal server banners. You can poll the status of the retrieval by getting the value of another new SNMP variable, sysConfigRadiusStatus.

Both variables now appear in the Ascend enterprise MIB, and are defined in the following manner:

```
sysConfigRadius OBJECT IDENTIFIER ::= { systemStatusGroup 6 }
sysConfigRadiusCmd OBJECT-TYPE
SYNTAX INTEGER {
    all( 1 )-- all configuration
    routes( 2 ),-- bridge/ip/ipx routes configuration
    pools( 3 ),-- ip address pool configuration
    nailed( 4 ),-- permanent and nailed configuration
    temrsrv( 5 ),-- term server banners
}
ACCESS read-write
STATUS mandatory
DESCRIPTION
"This variable instructs the system to re-read its RADIUS configura-
tion. This variable will return a general error if RADIUS is not con-
figured or if sysConfigRadiusStatus has a value of processing(2)."
::= { sysConfigRadius 1 }
sysConfigRadiusStatus OBJECT-TYPE
SYNTAX INTEGER {
    init( 1 ),-- configuration has not started
    processing( 2 ),-- configuration is in progress
    timeout( 3 ),-- configuration request timed-out
    error( 4 ),-- configuration received other error
    complete( 5 )-- configuration complete successfully
```



```

}
ACCESS read-only
STATUS mandatory
DESCRIPTION
    "This variable indicates the status of the last RADIUS configuration
    retrieval. This includes the boot time retrieval."
::= { sysConfigRadius 2 }

```

SNMP sysConfigTftpStatus object reports more states

In the Ascend MIB, the read-only sysConfigTftpStatus object in the System Status group previously only reported whether an SNMP-initiated download or upload passed or failed. It now reports a much wider variety of possible states.

```

sysConfigTftpStatus OBJECT-TYPE
SYNTAX INTEGER {
    ok( 1 ),-- tftp operation succeeded
    notFound( 2 ),-- file not found
    access( 3 ),-- access violation
    noSpace( 4 ),-- no disk space to write file
    badOp( 5 ),-- bad tftp operation
    exists( 7 ),-- file already exists
    noSuchUser( 8 ),-- no such user
    parameter( 9 ),-- parameter error
    busy( 10 ),-- tftp server cannot handle request
    noResources( 11 ),-- no memory for request
    timeout( 12 ),-- timed out
    unrecoverable( 13 ),-- unrecoverable error
    tooManyRetries( 14 ),-- too many retries
    createFile( 15 ),-- create file
    openFile( 16 ),-- open file
    inProgress( 17 )-- get/put request in progress
}
ACCESS read-only
STATUS mandatory
DESCRIPTION
    "This variable indicates the status of a save or restore operation
    through tftp."
::= { sysConfigTftp 2 }

```

SNMP now reports reasons for last reset

A new object in the systemStatusGroup in the Ascend MIB and a new trap report the reason for the last system reset. The possible errors are listed below.

sysLastRestartReason object

In the Ascend MIB, the read-only sysLastRestartReason object in the System Status group reports the reason the MAX reset.

```

sysLastRestartReason OBJECT-TYPE
    SYNTAX          INTEGER
    ACCESS           read-only
    STATUS           mandatory
    DESCRIPTION
        "The error code for the previous box restart. The error codes
        are identical to ones obtained via fatal-history from the
        debug monitor screen."
    ::= { systemStatusGroup 12 }

```

sysLastRestartReason trap

The sysLastRestartReason trap reports the reason the reason the MAX reset.

```

sysLastRestartReasonTrap TRAP-TYPE
    ENTERPRISE    ascend
    VARIABLES     { sysLastRestartReason }
    DESCRIPTION   "This trap is sent to all managers having the
                    alarm condition enabled if the
                    sysLastRestartReason is not unknown
                    (value of 0)."
```

::= 26

Definitions of fatal errors

This section describes the fatal errors that the MAX can report.

The following reset is caused when an Assert is placed in the code. This problem can be either hardware or software related. Contact Ascend Technical Support if you experience an FE1 reset.

```
FATAL_ASSERT = 1
```

The following resets are out-of-memory conditions, sometimes termed a memory leak.

```

FATAL_POOLS_NO_BUFFER = 2
FATAL_PROFILE_BAD = 3
FATAL_SWITCH_TYPE_BAD = 4
FATAL_LIF_FATAL = 5
FATAL_LCD_ERROR = 6
FATAL_ISAC_TIMEOUT = 7

```

The following reset is caused by a processor exception error.

```

FATAL_SCC_SPURIOUS_INT = 8
FATAL_EXEC_INVALID_SWITCH = 9

```

The following reset occurs because the MAX tried to allocate a mail message, and there were none left. A reset of this type is usually due to a memory leak.

```

FATAL_EXEC_NO_MAIL_DESC = 10
FATAL_EXEC_NO_MAIL_POOL = 11
FATAL_EXEC_NO_TASK = 12

```

FATAL_EXEC_NO_TIMER =	13
FATAL_EXEC_NO_TIMER_POOL =	14
FATAL_EXEC_WAIT_IN_CS =	15
FATAL_DSP_DEAD =	16
FATAL_DSP_PROTOCOL_ERROR =	17
FATAL_DSP_INTERNAL_ERROR =	18
FATAL_DSP_LOSS_OF_SYNC =	19
FATAL_DSP_UNUSED =	20
FATAL_DDD_DEAD =	21
FATAL_DDD_PROTOCOL_ERROR =	22
FATAL_X25_BUFFERS =	23
FATAL_X25_INIT =	24
FATAL_X25_STACK =	25
FATAL_ZERO_MEMALLOC =	27
FATAL_NEG_MEMALLOC =	28

The following reset is caused by a software loop.

FATAL_TASK_LOOP =	29
FATAL_MEMCPY_TOO_LARGE =	30
FATAL_MEMCPY_NO_MAGIC =	31
FATAL_MEMCPY_WRONG_MAGIC =	32
FATAL_MEMCPY_BAD_START =	33
FATAL_IDEC_TIMEOUT =	34
FATAL_EXEC_RESTRICTED =	35
FATAL_STACK_OVERFLOW =	36
FATAL_MBUF_PANIC =	38
FATAL_PROTECTION_FAULT =	40

The following entry is logged to the fatal-error table when the MAX has been manually reset, either in Diagnostic mode (with the RESET or NVRAMCLEAR commands), through the user interface, or through MIF.

Instead of a standard stack backtrace, the message includes the active security-profile index. The numbering is one-based, with 0 indicating an unknown security profile. On the MAX the Default profile is number 1, and the Full Access profile is number 9.

This reset is logged immediately before the MAX goes down.

FATAL_OPERATOR_RESET =	99
------------------------	----

As a complement to entry 99, the following entry is logged as the MAX is coming up. For a normal, manual reset, you should see a fatal error 99 followed by a fatal error 100.

FATAL_SYSTEM_UP =	100
-------------------	-----

Tunneling features

Layer 2 Tunnel Protocol (L2TP) support added

With this release, MAX units support Layer 2 Tunnel Protocol (L2TP), as specified in version 8 of the Internet Engineering Task Force (IETF) draft titled *Layer Two Tunneling Protocol “L2TP,”* dated November, 1997. L2TP enables you to connect to a private network by dialing into a local MAX, which creates and maintains an L2TP tunnel between itself and the private network.

Introduction

L2TP enables you to dial into a local ISP and connect to a private corporate network across the Internet. You dial into a local MAX, configured as an L2TP Access Concentrator (LAC), and establish a PPP connection. Attributes in your RADIUS user profile specify that the MAX, acting as an LAC, establish an L2TP tunnel. The LAC contacts the L2TP Network Server (LNS), which is connected to the private network. The LAC and the LNS establish an L2TP tunnel (via UDP), and any traffic your client sends is tunneled to the private network. Once the MAX units establish the tunnel, the client connection has a PPP connection with the LNS, and appears to be directly connected to the private network.

You can configure the MAX to act as either an LAC, an LNS, or both. The LAC performs the following functions:

- Establishes PPP connections with dial-in clients.
- Sends requests to LNS units requesting creation of tunnels.
- Encapsulates and forwards all traffic from clients to the LNS via the tunnel.
- De-encapsulates traffic received from an established tunnel, and forwards it to the client.
- Sends tunnel-disconnect requests to LNS units when clients disconnect.

The LNS performs the following functions:

- Responds to requests by LAC units for creation of tunnels.
- Encapsulates and forwards all traffic from the private network to clients via the tunnel.
- De-encapsulates traffic received from an established tunnel, and forwards it to the private network.
- Disconnects tunnels on the basis of requests from the LAC.
- Disconnects tunnels on the basis of expiration of the value you set for a user profile's Max-Connect-Time attribute. You can also manually disconnect tunnels from the LNS via SNMP, the terminal-server Kill command, or the DO Hangup command (which you access by pressing <control> D).

Note: With this release, a MAX acting as an LNS cannot send Incoming Call Requests to an LAC. Only an LAC can make requests for the creation of L2TP tunnels.

Configuring L2TP tunneling

This section describes how L2TP tunnels work between an LAC and an LNS. A client dials into an LAC, from either a modem or ISDN device, and the LAC establishes a cross-Internet IP connection to the LNS. The LAC then requests an L2TP tunnel via the IP connection.

The LNS is the terminating part of the tunnel, where most of the L2TP processing occurs. It communicates with the private network (the destination network for the dial-in clients) through a direct connection.

Figure 27 shows an ISP POP MAX, acting as an LAC, communicating across the WAN with a private network. Clients dial into the ISP POP and are forwarded across the Internet to the private network.

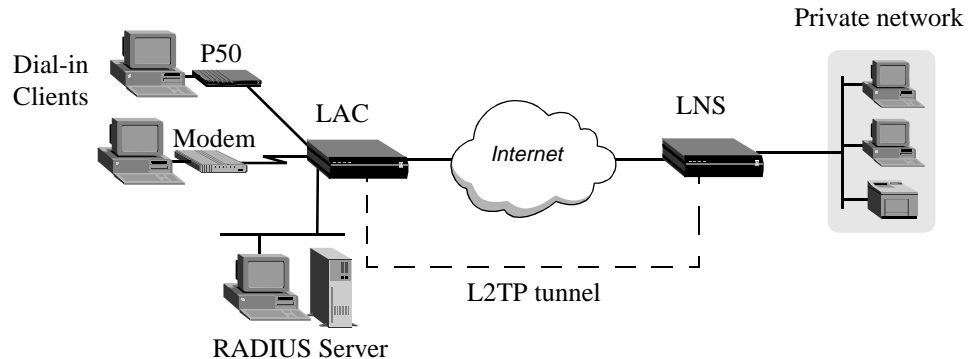


Figure 27. L2TP tunnel across the Internet

How the MAX creates L2TP tunnels

The dial-in client, the LAC, and the LNS establish, use and terminate an L2TP-tunnel connection as follows:

- 1 A client dials, over either a modem or ISDN connection, into the LAC.
- 2 On the basis of dialed number or after authentication (depending on the LAC configuration), the LAC communicates with the LNS to establish an IP connection.
- 3 Via the IP connection, the LAC and LNS establish a control channel.
- 4 The LAC sends an Inbound Call Request to the LNS.
- 5 Depending on the LNS configuration, the client might need to authenticate itself a second time.
- 6 After successful authentication, the tunnel is completed, and data traffic flows.
- 7 When the client disconnects from the LAC, the LAC sends a Call Disconnect Notify message to the LNS. The LAC and LNS disconnect the tunnel.

LAC and LNS mode

The MAX can function as an LAC, an LNS, or both. When configured as both, the MAX acts as an LAC on the basis of the dial-in client configuration. The MAX acts as an LNS when it receives an Inbound Call Request from an LAC.

Note: The MAX can support several simultaneous connections, some where it acts as an LAC, and some where it acts as an LNS. For any single connection, however, the MAX can operate as either an LAC or LNS, but not both.

Authentication

Either the LAC, the LNS, or both, can perform PAP or CHAP authentication of clients for which they create tunnels. If you configure the MAX to create tunnels on a per-line basis, only

Tunneling features

Layer 2 Tunnel Protocol (L2TP) support added

the LNS can perform authentication, because the MAX automatically builds a tunnel to the LNS for any call it that it receives on that line.

If you use RADIUS to configure L2TP on a per-user basis, and specify the Client-Port-DNIS attribute, the LAC does not perform PAP or CHAP authentication. When use specify Client-Port-DNIS, the tunnel is created as soon as the LAC receives the DNIS number and it matches a Client-Port-DNIS for any user profile. You can configure the LNS to perform PAP or CHAP authentication after the LAC and LNS establish the tunnel.

If you use RADIUS to configure L2TP, but do not specify the Client-Port-DNIS attribute, the LAC performs PAP or CHAP authentication before the tunnel is established. Once the tunnel is up, the LNS can perform authentication again on the client. Each client sends the same username and password during the authentication phase, so for each client, make sure you configure the LAC and LNS to look for the same usernames and passwords.

You can also direct the MAX to create an L2TP tunnel, from the terminal server, by using the `L2TP` command. You can configure authentication on the LNS, requiring users to authenticate themselves when they manually initiate L2TP tunnels from the terminal server.

Configuring the MAX as an LAC

The LAC is responsible for requesting L2TP tunnels to the LNS. You configure the LAC to determine when a dial-in connection should be tunneled, and you can specify the LNS used for the connection.

Understanding the L2TP LAC parameters

This section provides some background information about parameters used in configuring the MAX as an LAC:

Parameter	How it's used
L2TP enabled	To enable the MAX unit's LAC functionality, you must set L2TP to LAC or Both.
Line <i>n</i> tunnel type	Specifies whether the MAX should dedicate an entire WAN line to either L2TP or PPTP. If you want the MAX to establish tunnels on a connection-by-connection basis, set Line <i>n</i> tunnel type to None on all lines.
Route line <i>n</i>	Specifies the IP address of the LNS. This parameter applies <i>only</i> if you dedicate an entire WAN line to tunneling, using the Line <i>n</i> tunnel type parameter. If you want the MAX to establish tunnels on a connection-by-connection basis, leave Route line <i>n</i> blank for all lines.

Configuring the MAX as an LAC

To configure the MAX as an L2TP LAC, you must first enable L2TP LAC on the MAX, then configure how the MAX determines which connections are tunneled.

Configuring system-wide L2TP LAC parameters

To configure system-wide L2TP LAC parameters on the MAX:

- 1 Open the Ethernet > Mod Config > L2 Tunneling options menu.
- 2 Set L2TP Enabled to either LAC or Both.

Enabling L2TP tunneling for an entire WAN line

If you want the LAC to create L2TP tunnels for every call received on a specific WAN line:

- 1 Open the Ethernet > Mod Config > L2 Tunneling options menu.
- 2 For the line for which you are configuring LAC functionality (Line *n*), set Line *n* tunnel type to L2TP. For example, if you want to tunnel all calls received on the first WAN port (labelled WAN 1 on the MAX backpanel), set Line 1 tunnel type=L2TP.
- 3 Set Route line *n* to the IP address of the LNS.

Enabling L2TP tunneling on a per-user basis

You can configure RADIUS to direct the MAX to create L2TP tunnels for specific users. To do so, you use three standard RADIUS attributes: Tunnel-Type, Tunnel-Medium-Type, and Tunnel-Server-Endpoint. Table 28 describes them.

Table 28. RADIUS attributes for specifying L2TP tunnels

Attribute	Description	Possible values
Tunnel-Type (64)	Specifies which tunneling protocol to use for this connection.	PPTP or L2TP You must set this attribute to L2TP to direct the MAX to create an L2TP tunnel.
Tunnel-Medium-Type (65)	Specifies the protocol type, or medium, used for this connection. Currently, the MAX supports IP only. Future software releases will support additional medium types.	Currently, the only supported value is IP. You must set this attribute to IP.
Tunnel-Server-Endpoint (67)	Specifies the IP address or fully qualified host name of the LNS, if you set Tunnel-Type to L2TP, or PPTP Network Server (PNS), if you set Tunnel-Type to PPTP.	If a DNS server is available, you can specify the fully-qualified host name of the LNS. Otherwise, specify the IP address of the LNS in dotted decimal notation <i>n.n.n.n</i> , where <i>n</i> is a number from 0 to 255. You must set this attribute to an accessible IP host name or address.

Configuring the MAX as an LNS

When MAX acts as an LNS, it responds to requests by LAC units to establish tunnels. The LNS does not initiate outgoing requests for tunnels, so the configuration of MAX is simple. Proceed as follows:

- 1 Open the Ethernet > Mod Config > L2 Tunneling options menu.
- 2 Set L2TP Enabled to either LNS or Both.

New Parameters

The following new parameters support L2TP functionality:

L2TP Mode

Description: Specifies the system-wide type of L2TP functionality the MAX supports.

Usage: Specify one of the following values:

- LAC specifies that the MAX can function as an LAC only.
- LNS specifies that the MAX can function as an LNS only.
- Both specifies that the MAX can function as either an LAC or an LNS.
- None disables L2TP functionality on the MAX.
None is the default.

Example: L2TP Enable=LAC

Location: Ethernet > Mod Config > L2 Tunneling Options

See Also: Line *n* tunnel type, Route *n* line

Line *n* tunnel type

Description: Indicates whether the MAX should tunnel all calls received on the specified WAN line.

Usage: Specify one of the following values:

- L2TP directs the MAX to create L2TP tunnels for all calls received on the specified line.
- PPTP directs the MAX to create PPTP tunnels for all calls received on the specified line.
- None directs the MAX not to create tunnels on a per-line basis.
None is the default.

Example: Line 1 tunnel type=None

Dependencies: Line *n* tunnel type applies only if you set L2TP Mode to LAC or Both.

Location: Ethernet > Mod Config > L2 Tunneling Options

See Also: L2TP Mode, Route *n* line

Route line *n*

Description: Specifies the IP address of the L2TP Network Server (LNS) if you set Line *n* tunnel type to L2TP, or the IP address of the PPTP Network Server (PNS) if you set Line *n* tunnel type to PPTP.

Usage: Specify an IP address. The default is 0.0.0.0. If you accept the default, the MAX does not tunnel any call received on the WAN line specified in Line *n* tunnel type.

Example: Route Line 1=10.10.10.10

Dependencies: Route line *n* applies only if you set L2TP Mode to LAC or Both.

Location: Ethernet > Mod Config > L2 Tunneling Options

See Also: L2TP Mode, Line *n* tunnel type

RADIUS attributes

The following RADIUS attributes support L2TP and PPTP tunneling:

Tunnel-Type (64)

Description: Specifies the type of tunneling protocol to create.

Usage: You can specify the following values for Tunnel-Type:

- PPTP
- L2TP

Example: Tunnel-Type=L2TP

Dependencies: For the MAX to correctly create an L2TP tunnel, you must set Tunnel-Medium-Type to IP and set Tunnel-Server-Endpoint to the IP address of an accessible LNS, in addition to setting Tunnel-Type to L2TP.

For the MAX to correctly create an PPTP tunnel, you must set Tunnel-Medium-Type to IP and set Tunnel-Server-Endpoint to the IP address of an accessible PNS, in addition to setting Tunnel-Type to PPTP.

See Also: Tunnel-Medium-Type (65), Tunnel-Server-Endpoint (67)

Tunnel-Medium-Type (65)

Description: Specifies the type of medium supported over the tunnel specified in the Tunnel-Type attribute.

Currently, the MAX supports only the IP medium type.

Usage: Currently, you can specify IP for the value of Tunnel-Medium-Type.

Example: Tunnel-Medium-Type=IP

Dependencies: For the MAX to correctly create an L2TP tunnel, you must set Tunnel-Type to L2TP and set Tunnel-Server-Endpoint to the IP address of an accessible LNS, in addition to setting Tunnel-Medium-Type to IP.

For the MAX to correctly create a PPTP tunnel, you must set Tunnel-Type to PPTP and set Tunnel-Server-Endpoint to the IP address of an accessible PPTP Network Server (PNS), in addition to setting Tunnel-Medium-Type to IP.

See Also: Tunnel-Type (64), Tunnel-Server-Endpoint (67)

Tunnel-Server-Endpoint (67)

Description: Specifies the fully-qualified host name or IP address of the network server to contact for building a tunnel. If you set Tunnel-Type to L2TP, Tunnel-Server-Endpoint indicates the IP address of the LNS. If you set Tunnel-Type to PPTP, Tunnel-Sever-Endpoint indicates the IP address of the PNS.

Usage: Specify the primary home agent in the following format:

Tunnel-Server-Endpoint=*"hostname | ip_address"*

where:

- *hostname* is the fully-host name of the network server.
- *ip_address* is the network server's IP address in dotted decimal notation.
Specify an IP address if the network server does not have access to a DNS server.

You can specify a host name or IP address, but not both.

Example: To specify the network server maxSF.home.com at IP address 10.10.10.10, specify one of the following lines in the RADIUS user profile:

Tunnel-Server-Endpoint=10.10.10.10

Tunnel-Server-Endpoint=maxSF.home.com

Dependencies: For the MAX to correctly create an L2TP tunnel, you must set Tunnel-Type to L2TP and Tunnel-Medium-Type to IP, in addition to specifying the IP address of an accessible LNS.

For the MAX to correctly create an PPTP tunnel, you must set Tunnel-Type to PPTP and Tunnel-Medium-Type to IP, in addition to specifying the IP address of an accessible PPTP Network Server (PNS).

See Also: Tunnel-Type (64), Tunnel-Medium-Type (65)

Maximum number of ATMP Tunnel sessions can be set

You can now specify the maximum number of active ATMP tunnel sessions allowed through a unit configured as an ATMP Home Agent. Previously, there was no way to limit ATMP sessions. With ATMP enabled, it was possible for all active sessions to be ATMP sessions.

Max ATMP Tunnels

Description: Defines the maximum number of active ATMP sessions for units configured as an ATMP Home agent.

Changes take effect after the Connection Profile is saved, the connection is cleared, then reestablished.

Usage: Press Enter to open the text field. Type the number of simultaneous ATMP sessions you want to allow through this ATMP Gateway. The default, 0 (zero), disables the parameter.

Dependencies: Applies only to units configured as ATMP Home agents.

Location: Ethernet > Connections > *any profile* > Session Options menu.

See Also: ATMP Mode, ATMP Gateway.

ATMP inactivity timer

You can now configure a timer for ATMP tunnels, indicating the number of minutes the Ascend unit maintains an idle tunnel before disconnecting it.

Overview

ATMP tunnels between an ATMP foreign agent and an ATMP home agent are disconnected when the foreign agent detects the client, for which the tunnel was created, disconnects. However, if you reset a foreign agent, the home agents of any existing tunnels get no notification that the tunnels should be disconnected.

When an Ascend unit, acting as an ATMP foreign agent, restarts, tunnels that were established to any home agent are not normally cleared, because the home agent is never informed that the mobile clients are no longer connected. The unused tunnels continue to consume resources on the home agent until the Ascend unit acting as the home agent, is restarted. To enable the home agent to reclaim the resources held by unused tunnels, you can configure the Idle limit parameter to indicate how long a home agent maintains an idle tunnel before disconnecting it.

If you set the ATMP inactivity timer, it will not affect any other idle timers you might have configured on the MAX. The MAX can apply more than one idle timer to a particular connection.

Idle limit

Description: Specifies the number of minutes the Ascend unit, configured as an ATMP home agent, maintains an idle ATMP tunnel before it disconnects the tunnel.

Changes you make to Idle limit are enabled for any new tunnels. Existing tunnels are not affected.

Usage: Specify the number of minutes the home agent should maintain any ATMP tunnel before disconnecting it.

Disable the ATMP inactivity timer by setting it to 0. This is the default.

Example: `Idle limit=15`

Dependencies: Idle limit is not applicable if you set the Ethernet > Mod Config > ATMP > ATMP Mode parameter to either Disabled or Foreign.

Location: Ethernet > Mod Config > ATMP

See Also: ATMP Mode, Type, Passwd, SAP Reply, UDP Port

Administration features

Support for Finger remote user information protocol (RFC 1288)

The MAX now supports Finger remote user information protocol (RFC 1288).

Overview

You can use Finger to get information about users currently logged into the MAX. This includes the host address, name, port, and channel. Note that for security reasons the MAX does not forward Finger requests. Refer to RFC 1288 for complete details of the Finger protocol.

Using the Finger command

To use the Finger command to get information about a particular use, type:

```
finger jane-doe@a.b.c.d
```

I Session	Line:	Slot:	Data	Service	Host	User
O ID	Channel	Port	Rate	Type[mpID]	Address	Name
O 214933581	1:2	9:1	56K	MPP[1]	192.168.4.9	jane-doe

Where *jane-doe* is the name of a user currently logged into the MAX and *a.b.c.d* is the IP address of the MAX. The user name must match the name in a Connection profile, Name/Password profile, or RADIUS User profile.

To use the Finger command to get information about all the users currently logged into the MAX, type:

```
finger @a.b.c.d
```

Where *a.b.c.d* is the IP address of the MAX. The user name must match the name in a Connection profile, Name/Password profile, or RADIUS User profile.

I Session	Line:	Slot:	Data	Service	Host	User
O ID	Channel	Port	Rate	Type[mpID]	Address	Name
O 214933581	1:2	9:1	56K	MPP[1]	192.168.4.9	arwp50
O 214933582	1:6	9:2	56K	MPP[1]	MPP Bundle	arwp50
I 214933583	1:1	3:1	28800	Termsrv	N/A	trm-havnor

Configuring Finger

To enable the MAX to respond to Finger requests:

- 1 Open Ethernet > Mod Config.
- 2 Set Finger to Yes.

Error messages

The Finger command may return the following messages:

Table 29. Finger messages

Message	Explanation
No Active Users	No users are currently logged into the MAX.
Finger online user list denied.	Finger on this MAX is disabled.
Finger forwarding service denied.	The MAX does not forward Finger requests.

Parameter reference

Finger

Description: Enables or disables the Finger remote user information protocol (RFC 1288). Finger returns information about users currently logged into the MAX. Note that for security reasons the MAX does not forward Finger requests.

Usage: Specify one of the following values:

- Yes enables the Finger protocol.
- No disables the Finger protocol.

Location: Ethernet > Mod Config

Terminal server show users command added

A terminal server command has been added that displays a list of user sessions active on a system. Each user session is identified by the sessionID, with additional information about the session. The show users command has also been added to the online help for the show command.

How the show users command works

To display a list of active user session on an Ascend MAX or Pipeline 400, type

show users

at the terminal server prompt. The sessions are listed with one session to each line. If there are no active sessions the message "No active connections" is displayed.

Administration features

Terminal server show users command added

The following table shows the information the show users command will return.

Note: This feature has only been tested for terminal server, PPP, MP, and MPP sessions and a Net/T1 card. Display of data from other session types or interfaces is indeterminate.

IO	O (outgoing call) I (incoming call)
Session ID	Session identifier for the session. This number can be cross- referenced to other system functions that also use the sessionID. It is the same as Acct-Session-ID in RADIUS.
Line Channel	The WAN line and channel associated with the session. If the session is over a frame relay connection, the line and channel of the underlying frame relay connection are listed. For connections with multiple channels, the channel number is replaced with a dash (-).
Slot Port	The slot and port of the service being used by the session. This can be: <ul style="list-style-type: none">the number of a slot containing a modem card and the modem on that cardthe virtual slot of the MAX's bridge/router, with port giving the virtual interfaces to MAX's bridge/router starting with 1 for the first session of a multichannel session
Data Rate	The bearer capacity or modem speed as appropriate to the session type. For multi-channel calls (not including MP family calls where the different channels are associated with distinct calls) this is the bandwidth of the entire call, which may be shared with other sessions.
Service Type	The type of session. This can be Termsrv or a protocol name. For MP and MPP sessions the ID of the session is listed following the encapsulation protocol name to allow identification of the MP bundle to which a session belongs. The special values Initial and Login document the progress of a session: <ul style="list-style-type: none">Initial identifies sessions that do not yet have a protocol assigned.Login identifies sessions waiting at the terminal server login prompt.

Host Address	The network address of the remote host originating the session. For some sessions this field is N/A. If the session is part of an outgoing MPP bundle, this field is set to "MPP Bundle" for the secondary members of the bundle.
User Name	The name associated with the session, usually the name of the remote host. For terminal server sessions this is the login name. Prior to login completion this field will show the string "modem x:y" where x and y are the slot and port of them modem servicing the session.

On-line help for the Show Users command

The Show Users command has been added to the on-line Help. To see a list of Show commands, type

Show ?

Example data displayed by show users

In the first example output, two sessions running at 56K are indicated. The first session came in on line 5 and the second on line 6.

- Both are part of MP bundle 1
- Both are incoming ISDN calls from arwp50.

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

```
ascend% show users
```

```
I Session  Line:  Slot: Data Service  Host      User
O ID      Channel Port Rate Type[mpID] Address   Name
I 219846946 1:5    9:1  56K  MPP[1]   192.168.6.145 arwp50
```

The second example shows seven sessions:

- The first two sessions are part of an outgoing MP bundle (bundle 2) and are on lines 1 and 2 of the first T1 line.
- The third session is an outgoing MPP bundle consisting of one call on line 3.

- Line 4 is a single analog call to the terminal server using the login name trmhavnor. This call is connected at 1200 B over modem 1 in the modem card in slot 3.
- Line 5 is an analog PPP call connected at 28800 over modem 2. The second T1 line is connected to another MAX (Iffish) and is configured for a twelve-channel nailed frame relay connection using the profile "dialslip" connected at 28800 through the 12th modem.

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

```
ascend% show users
```

I Session	Line:	Slot:	Data	Service	Host	User
O ID	Channel	Port	Rate	Type[mpID]	Address	Name
O 219846948	1:1	9:1	56K	MPP[2]	192.168.6.145	arwp50
O 219846949	1:2	9:2	56K	MPP[2]	MP Bundle	arwp50
O 219846950	1:3	9:3	56K	MPP[3]	192.168.10.1	roke-gw

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 OK u=Torning c=185 p=60 s=64000 r=64000
... ASCEND: slot 9 port 1, line 1, channel 1, Call Disconnected
... ASCEND: call 1 CL OK c=20 p=40 s=31200 r=33600
... ASCEND: call 3 AN slot 3 port 2 VOICE
... ASCEND: call 3 CL OK c=185 p=31
... ASCEND: slot 3 port 2, line 1, channel 2, Call Disconnected
```

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.

Receive and transmit rates now reported

The terminal server show users command, RADIUS, and SNMP now report and display the receive and transmit data rates for user sessions. Previously, the transmit rate was not recorded or reported separately from the receive rate.

Displaying user session statistics

At the terminal server prompt type the following to display user session status:

show user

The new show users format is:

I	Session	Line:	Slot:	Tx	Rx	Service	Host	User
O	ID	Chan	Port	Data	Rate	Type[mpID]	Address	Name
O	231849873	1:1	1:9	56K	56K	MPP[1]	192.168.68.2	jdoe
I	231849874	1:1	3:1	28800	33600	Termsrv	N/A	Modem
	3:1							

at the terminal server prompt.

Previously, the Data Rate column displayed the bearer capacity or modem speed as appropriate to the session type. The Data Rate column has been removed from the display, and two new columns have been added:

- Rx (Receive) Rate
Shows the receive data rate in bits per second.
- Tx (Transmit) Rate
Shows the transmit data rate in bits per second.

Note: This feature does not apply to LAN modem cards. These cards do not support any asymmetric data rate connection types, so the transmit rate is the same as the receive rate.

New RADIUS attributes for receive and transmit rates

The RADIUS dictionary has been changed to accommodate the new session status display. Ascend-Xmit-Rate (255) has been added to the RADIUS dictionary to specify the transmit baud rate. Ascend-Data-Rate (197) previously was described as reporting the data rate of the connection in bits per second. This attribute now should be described as reporting the receive baud rate.

Ascend-Xmit-Rate (Attribute 255)

Description: Specifies the transmit baud rate for the connection.

Dependencies: The Ascend-Xmit-Rate attribute is sent in Accounting-Request packets at the end of a session under these conditions:

- The Accounting-Request packet has Acct-Status-Type=Stop.
- The Auth parameter is set to a value other than RADIUS/LOGOUT.

The attribute is still sent with the Accounting-Request packet whether the connection is authenticated or not.

Changes to the Ascend MIB

You can use SNMP to obtain the same received and transmit data rate statistics as shown by the show users command. To accommodate this, three new objects have been added to the Ascend MIB, and three objects have changed their meaning:

- callStatusXmitRate (callStatusEntry 14) (1.3.6.1.4.1.529.11.2.1.14)
callStatusXmitRate is similar to callStatusDataRate, but provides the transmit data rate for an ISDN call or the baud rate for a modem call. callStatusDataRate now indicates the receive data rate for an ISDN call or the baud rate for a modem call.
- callActiveXmitRate (callActiveEntry 14) (1.3.6.1.4.1.529.11.16.1.14)
The callActiveTable contains a list of active call status entries. callActiveXmitRRate is a entry in callActiveEntry, which contains object variables that describe an active call's status. callXmitRate provides the transmit data rate for an ISDN call or the baud rate for a modem call.
- eventXmitRate (eventEntry 23) (1.3.6.1.4.1.529.10.4.1.23)
eventXmitRate is similar to eventDataRate, an object variable in an eventEntry object in the eventTable. eventXmitRate provides the transmit data rate for an ISDN call or the baud rate for a modem call.

Syslog enhancements

Syslog has been enhanced to report more information about calls. These enhancements include detailed information about authenticated calls when they disconnect and the Dialed Number Identification Service (DNIS) and Calling Line ID (CLID) for each call.

Call information forwarded to syslog when call terminates

The MAX now sends the syslog daemon information about an authenticated connection when the call terminates. To enable this feature, a new parameter, LogCallInfo, has been added to the Ethernet > Mod Config > Log submenu.

When LogCallInfo is set to EndOfCall, the MAX sends a one-line syslog message to the syslog host when an authenticated call terminates. The message contains the following information:

- Connection information
 - station-name
 - calling phone number
 - called phone number
 - encapsulation protocol
 - datarate (in bps)
 - progress code/disconnect reason
- Authentication information

- time spent before authenticating (in seconds)
- bytes/packets received during authentication
- bytes/packets sent during authentication session
- Session information
 - time spent in session (in seconds)
 - bytes/packets received during session
 - bytes/packets sent during session

For example:

```
"Conn=( "cvonk-p50" 5105558581->? PPP 56000 60/185) \  
Auth=(3 347/12 332/13) \  
Sess=(1 643/18 644/19), Terminated"
```

If some of the information is not available, that field is displayed as either a question-mark (for strings) or a zero (for numerals).

Note: This feature is intended for diagnostic support. It uses User Datagram Protocol (UDP), which provides no guaranteed delivery, so it should not be used for billing purposes.

Parameter reference

This section describes the new parameter to support this feature.

LogCallInfo

Description: Specifies whether the MAX should send connection, authentication and session information about authenticated calls to Syslog when the call terminates. The MAX reports the following information:

- Connection information
 - station-name
 - calling phone number
 - called phone number
 - encapsulation protocol
 - datarate (in bps)
 - progress code/disconnect reason
- Authentication information
 - time spent before authenticating (in seconds)
 - bytes/packets received during authentication
 - bytes/packets sent during authentication session
- Session information
 - time spent in session (in seconds)
 - bytes/packets received during session

Administration features

Syslog messages generated when Security profile activated

- bytes/packets sent during session

Usage: Specify one of the following values. The Default is None:

- None specifies that the MAX does not report any information about authenticated calls to Syslog when the call disconnects.
- EndOfCall specifies that the MAX reports connection, authentication and session information about authenticated calls when they disconnect.

Dependencies: Keep in mind the following additional information:

- LogCallInfo does not apply if Syslog is not enabled.

Location: Ethernet > Mod Config > Log

See Also: Syslog, Log Host, Log Port, Log Facility

Syslog now reports DNIS and CLID information

The MAX did not previously report Calling Line ID (CLID) and Dialed Number Identification Service (DNIS) to Syslog.

Syslog messages that were previously displayed like this:

```
[1/1/1/5] [MBID 2] Incoming Call
```

will now appear as follows, whenever possible:

```
[1/1/1/5] [MBID 2; 5107891212->7242] Incoming Call
```

This format indicates that the Caller ID (5107891212 in the example above) is calling the number 7242.

Syslog messages generated when Security profile activated

This feature enables you to detect and handle unauthorized Telnet or serial-port sessions with the MAX. When a user activates a Security profile, the MAX generates a Syslog message notifying you that the event occurred.

New Syslog messages

A user can activate a Security profile in a Telnet session or a serial-line COM port session by selecting the Security profile and specifying the proper password. When a user activates a Security profile, the new Syslog messages show the name of the Security profile, the IP address of the Telnet client or the COM port number, and the local IP address.

The EventSyslog message has one of these formats:

```
^DP(assword)ASCEND: "<profile_name>" ... for <remote_IP> on <local_IP>
ASCEND: "<profile_name>" ... from <COM_port> on <local_IP>
```

- The <profile_name> argument specifies the name of the activated Security profile.

- The <remote_IP> argument specifies the IP address of the Telnet client.
- The <local_IP> argument specifies the local IP address of the MAX.
- The <COM_port> argument specifies the COM port number for the session.

On system login, the MAX does not generate a Syslog message for the Default Security profile; for all events other than system login, the MAX generates a Syslog message for the Default Security profile. If Syslog is enabled, messages at LEVEL_NOTICE appear when a user activates a Security profile and the MAX accepts the Security profile password.

These two messages signal that a Telnet client has enabled a Security profile:

```
Jan 10 10:05:17 eng-lab-141 ASCEND: "Full Access" security profile
enabled for 206.65.212.9 on 192.168.6.141.
```

```
Jan 10 10:07:26 eng-lab-141 ASCEND: "Default" security profile enabled
for 206.65.212.23 on 192.168.6.141.
```

This message signals that a COM port user has enabled the Full Access profile:

```
Jan 10 10:03:52 eng-lab-141 ASCEND: "Full Access" security profile
enabled from com port 0 on 192.168.6.141.
```

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.

Defender authentication enhancements

Syslog messages are now logged when a telnet client logs in and when a Security profile is activated. This feature helps detect and control unauthorized telnet sessions and possible security breaches in the MAX.

When a telnet "accept" takes place, a syslog message now is logged showing the IP address of the telnet client. When a caller again tries to make a connection, the router begins the process of connecting to hosts until it either succeeds or fails after trying one authentication host after another until the entire list of hosts has been processed. This allows the router to authenticate a subsequent user when an authentication host becomes available.

Alternate authentication host configuration

To implement this change, the Defender authentication method now supports addresses for more than one authentication host. In previous releases, the Auth Host #2 and #3 parameters were N/A when Auth was set to Defender. In the current release, the Defender configuration could include an alternate authentication host; for example:

```
Mod Config
Auth...
Auth Host #1=137.175.80.62
Auth Host #2=0.0.0.0
Auth Host #3=137.175.80.24
Auth Port=2626
Auth Src Port=0
Auth Timeout=30
Auth Key=
```

There is no problem with "skipping" an authentication host, such as the null address shown for Auth Host #2.

Syslog messages

This release also supports a set of syslog messages reporting the status of the Defender authentication subsystem. The new syslog messages are reported with "LOG_DEFAULT" and "LEVEL_INFO" priorities; for example:

```
Nov 14 15:59:34 137.175.85.20 ASCEND: AuthHost 137.175.81.24 Acti-
vated
Nov 14 15:51:10 137.175.85.20 ASCEND: AuthHost 137.175.81.24 Fails
auth
Nov 14 15:51:10 137.175.85.20 ASCEND: AuthHost 137.175.80.24
Refuses connect
Nov 14 16:03:05 137.175.85.20 ASCEND: AuthHost 137.175.81.24 Closed
connection
Nov 14 16:05:59 137.175.85.20 ASCEND: AuthHost 137.175.81.24
Address Changed
Nov 14 16:06:31 137.175.85.20 ASCEND: AuthHost 137.175.81.24 New
Authmethod
```

All Defender syslog messages report the standard Ascend header plus Defender-formatted detail: "AuthHost xx.yy.zz.aa StatusX" where "StatusX" has the following identities and meanings:

- **Activated**
A Defender Host has been found, with a successful connection established. This state is reported when an authentication session is active and ready to authenticate.
- **Fails auth**
A Defender Host has been found, but the router and the Defender authentication server do not agree on their mutual authentication key.
- **Refuses connect**
The host either is not responding at all, or has no active Defender Server running.
- **Closed connection**
An active Defender authentication server has ended its connection. This would reflect either a failure of the server software or explicit request for the server to shutdown by an administrator.
- **Address Changed**
The router administrator has elected to change the IP address of the active authentication server, its port number, or the authentication key. This forces the Defender authentication subsystem to close its connection with the active server and start searching for a new one.
- **New Authmethod**
The router administrator has changed the authentication method from "DEFENDER" to something else, causing the Defender authentication subsystem to break an active connection.

RLOGIN enhanced -l option

The "-l" option of the rlogin feature has been modified. Rlogin functionality has not been changed. However, the "-l" option can now be specified either before or after the hostname.

Starting an Rlogin session

From the command line, the user can enter the RLOGIN command in either of the following formats:

- `rlogin [-e<char> -l<username>] <host-name>`
- `rlogin [-e<char>] <host-name> [-l<username>]`

If DNS is configured in the MAX Ethernet Profile, you can specify the host's name, for example:

```
rlogin myhost
```

If DNS is not configured in the Ethernet Profile, you must type the host's IP address, for example:

```
rlogin 10.2.2.2
```

The -e option is used to set the escape character to the specified character. The default for the escape character is a tilde (~).

The -l option is used to specify a login name other than the name used to log into the terminal server. However, the -l option does not apply to logins authenticated by RADIUS.

To close the Rlogin session, the user must type a carriage return followed by the escape character. For example:

<CR>~

TFTP checks for compatibility of downloaded files

With this release, the Ascend unit compares the software to be TFTP-uploaded to the currently loaded software. If the platform or network interface does not match, the Ascend unit aborts the upload and displays information about why the abort occurred. The MAX 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 Ascend unit. Previously, you were able to upload any software to any Ascend unit. 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 Ascend unit 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 Ascend unit upload the older software without performing the compatibility check.

Examples

In the following example, a user attempts to use TFTP to upload a MAX 1800 software load (b.m18) to a MAX 4000 running t.m40:

- 1 From the vt100 interface, accesses the diagnostics monitor.
- 2 Enters the following command:
- 3 The MAX 4000 displays the following information to the screen:

```
tloadcode tftpserver.ascend.com b.m18

saving config to flash

.....
.
loading code from tftpserver.ascend.com
file /tftpboot/b.m18...
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 MAX 4000 has compared the uploading file, b.m18 to its currently-loaded file, t.m40. 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 MAX 4000 that uses this feature:

- 1 From the vt100 interface, accesses the diagnostics monitor.
- 2 Enters the following command:
- 3 The MAX 4000 displays the following information to the screen:

```
tloadcode tftpserver.ascend.com t.m140

saving config to flash

.....
.
loading code from tftpserver.ascend.com
file /tftpboot/t.m40...
thin load:
This load has no platform identifier. Proceed with caution.
Download aborted. Use 'tloadcode -f' to force.
```

In the previous example, the user decides that he or she requires the older version and forces the upload. The following messages are displayed:

- 1 User enters the following command
- 2 The MAX 4000 displays the following messages:

```
tloadcode -f tftpserver.ascend.com t.m140

Download forced by user...
.....
.....
.....
>
```

Cexample now uses a session server key

Cexample is a sample program included in the Ascend RADIUS release. It demonstrates how to send commands to the RADIUS server for changing filters and terminating sessions. Each RADIUS request must have an attribute that ties the command to a particular session. One such attribute is Ascend-Session-Svr-Key (151). This release adds support for this attribute to the Cexample program.

New command-line option

The Cexample command line now includes a new `-k sessKey` option:

```
cexample -H servAddr -P port -K secret [-s sessID][-u name][-i ipAddr][-k sessKey][-x]
```

The `sessKey` argument specifies the value of the Ascend-Session-Svr-Key. This attribute enables the MAX to match a user session with a client request to perform certain operations, such as disconnecting a session or changing a session's filters. The client sends Ascend-Session-Svr-Key to the RADIUS server in a Disconnect-Request or Change-Filter-Request packet when it initiates an operation. In addition, Ascend-Session-Svr-Key appears in a RADIUS Accounting-Start packet when a session starts.

Modem features

T1-CSU Status window added to the MAX 2000

You must enter the value in two-digit hexadecimal format. For example, enter the number "2" as "02", and the letter "a" as "0a". For example, if the session server key is 0ab34578efbc8423, enter this command line

```
cexample -H srl -P 1700 -K secret -k 0ab34578efbc8423
```

T1-CSU Status window added to the MAX 2000

The MAX 2000 now includes a reference to the T1-CSU Status window. Previously the Status window for the MAX 2000 leased T1 line was difficult to find.

There is now a pointer to the MAX 2000 leased T1 line status window. The screen below illustrates the change to the user interface:

```
Main Status Menu
00-000 System
10-000 Net/T1
20-000 Empty
30-000 Empty
40-000 See 10-200
50-000 Ethernet
```

This indicates that to view the T1-CSU status, select 10-000 Net/T1 > 10-200 Line Errors. For information on the MAX T1 line errors, refer to the documentation that came with your MAX.

Modem features

Rockwell code version

Rockwell K56 code version 1.160t is now supported on the MAX Series56 digital modems.

New 56K digital modem cards

Three new 56Kbps Rockwell modem cards are now supported on all MAX units except the MAX 200Plus:

- K56 Modem-8DIG. 8 digital modems (MAX 1800 only)
- K56 Modem-12DIG. 12 digital modems (MAX 2000, MAX 4000, MAX 4002, MAX 4004)
- K56 Modem-16DIG. 16 digital modems (MAX 1800, MAX 4000, MAX 4002 and MAX 4004 only)

Note that you cannot install the new 56Kbps modem cards into the same chassis with existing modem cards. In addition, all 56Kbps modem cards in a MAX must be of the same density.

To provide better performance, these new parallel-mode digital modems contain an onboard Digital Signal Processor (DSP). These modems support the following standards:

- K56Flex
- V.34
- V.32
- V.FC

User interface changes

When you install one of the new digital modems into a MAX, the main edit menu in the VT-100 interface displays new menus indicating the type of modem card installed:

```
Main Edit Menu
  00-000 System
  10-000 Net/T1
  30-000 Empty
  40-000 K56 Modem-8DIG
  50-000 K56 Modem-12DIG
  60-000 K56 Modem-16DIG
  70-000 Empty
  80-000 Empty
  90-000 Ethernet
  A0-000 Ether Data
  B0-000 Serial WAN
```

In addition, the Modem diagnostic screen is expanded to include 16 modems.

```
40-200 Modem Diag
Modem #1= enable modem
Modem #2= enable modem
Modem #3= enable modem
Modem #4= enable modem
Modem #5= enable modem
Modem #6= enable modem
Modem #7= enable modem
Modem #8= enable modem
Modem #9= enable modem
Modem #10= enable modem
Modem #11= enable modem
Modem #12= enable modem
Modem #13= enable modem
Modem #14= enable modem
Modem #15= enable modem
Modem #16= enable modem
```

Disable modem capability added for MAX 1800 and MAX 2000

You may now enable, disable modem capability for MAX 1800 and MAX 2000. For MAX 2000, you may also disable and channel modem capability.

The new settings are at the following locations:

- In the VT100 interface; Mod Config>Module Name>Modem Slot>*any modem #*
- In the SNMP interface; in misc.mib, slotMdmTable.

New Mod Config capability

You may set any modem that is operated by your MAX. There are two settings that you may choose from:

- enable modem, *modem ready for configuration and use*
- disable modem, *modem cannot be configured or used*

There is an additional setting for the MAX 2000 T1 interface:

- disable + channel, *disable and channel*

New values for SNMP variables

You may set the slotMdmItemConfig in slotMdmTable to configure modem usage. Here are the values you may set MAX:

2, enable

3, disable

This value may only be set for MAX 2000 T1 interface.

4, disable + channel

You may monitor modem resources by using the lanModemGroup.

56k Modem Numbering

Previously, when using the Show Modems Terminal Server command, modems on an 8-MOD modem card were numbered 1 to 8. Modems on a 12-MOD modem card were numbered 1 to 12. Modem cards that now support 56k technology, are not numbered sequentially. This numbering does not affect functionality.

8-MOD modem numbering

Modems in the 8-MOD modem card are numbered 0, 1, 2, 3, 6, 7, 10, 11.

For example, if you have an 8-MOD modem card in slot 8 in a MAX 4000, the Show Modems command in the Terminal Server displays the following output:

```
ascend% show modems
slot:item    modem    status
8:0          1        idle
8:1          2        idle
8:2          3        idle
```

8:3	4	idle
8:6	5	idle
8:7	6	idle
8:10	7	idle
8:11	8	idle

12-MOD modem numbering

Modems in the 12-MOD modem card are numbered 0, 1, 2, 3, 4, 5, 6, 7, 8, 9, 12, 13

For example, if you have an 12-MOD modem card in slot 8 in a MAX 4000, the Show Modems command in the Terminal Server displays the following output:

```
ascend% show modems
slot:item    modem    status
8:0          1        idle
8:1          2        idle
8:2          3        idle
8:3          4        idle
8:4          5        idle
8:5          6        idle
8:6          7        idle
8:7          8        idle
8:8          9        idle
8:9          10       idle
8:12         11       idle
8:13         12       idle
```

Ability to select between V.34 and K56Flex modulation

MDM Modulation is a new parameter that provides the ability to configure K56Flex-capable MAX modems to either negotiate for a maximum baud rate of 56K or a maximum baud rate of 33.6K. By default, all MAX modems attempt to negotiate for the maximum baud rate possible. This parameter enables you to restrict K56Flex modems from negotiating above V.34 modulation rates.

MDM Modulation

Description: Specifies the maximum modulation rate negotiated by K56Flex-capable MAX modems.

The selection specifies the maximum baud rate the MAX modem negotiates. The modems will continue to negotiate to slower baud rates based on line conditions or restrictions of the far end modem.

Usage: Press Enter to cycle through the choices.

- K56 instructs the MAX to negotiate using the K56Flex protocol.
This is the default.
- V.34 instructs the MAX not to negotiate using the K56Flex protocol.

Dependencies: If TS Enabled=No, the MDM Modulation does not apply (MDM Modulation=N/A).

Location: Ethernet > Mod Config > Tserv Options

X.25 features

X.25 T3POS support

MAX units with X.25 now support the T3POS protocol, which can be used to send point of sale (POS) data over the ISDN D channel.

This feature provides X25 Transaction Processing Protocol for Point-of-Service (T3POS) support for the MAX over the existing Ascend X.25 stack. T3POS is a character-oriented, frame-formatted protocol designed for point-of-service (POS) transactions through an X.25-based packet switched network. T3POS allows you to send data over the ISDN D channel while continuing to send traffic over both B channels. The T3POS protocol involves three parties: the T3POS DTE (DTE for short), the T3POS PAD (PAD for short) and the T3POS Host (host for short). See the following figure.




Figure 30. T3POS set up

A typical use of T3POS is to perform credit card authorization over the D channel while using the B channels to transmit inventory control data and other traffic. An example T3POS setup is illustrated in the following figure.

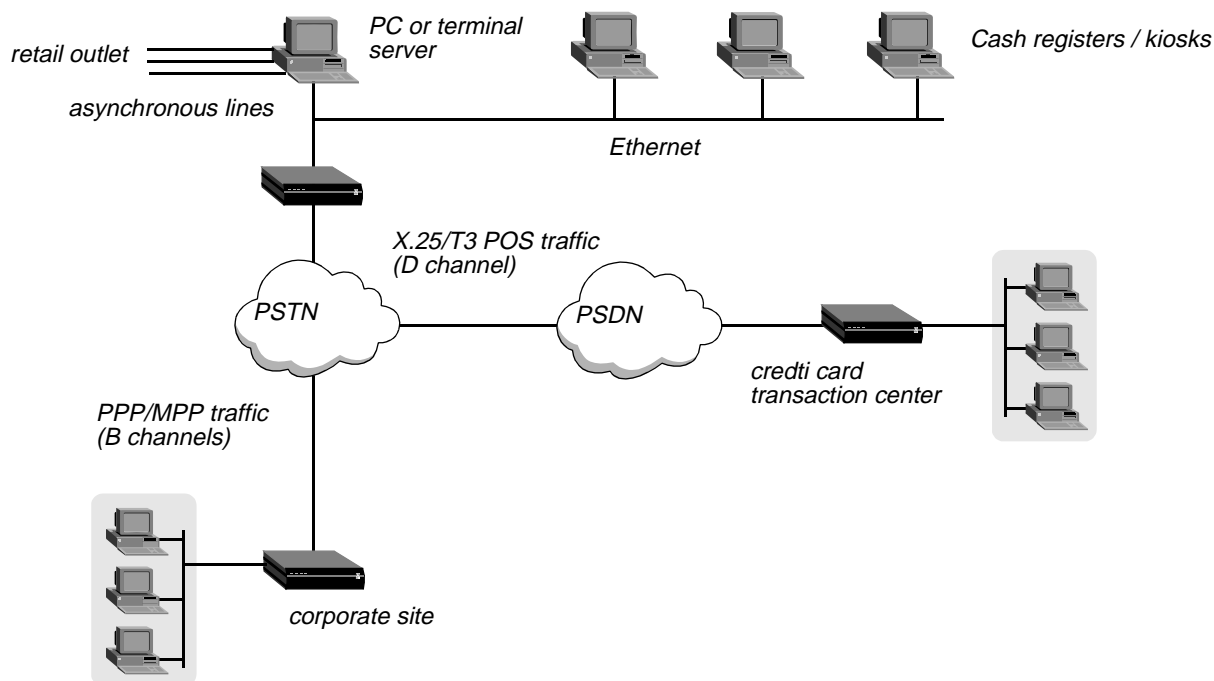


Figure 31. Example T3POS configuration

The Ascend T3POS implementation supports the following T3POS features:

- Local, Transparent, Blind and Binary-local mode
- T1-T6 timers
- All the control characters as described in Bellcore GR-2803
- Error recovery procedures as described in Bellcore GR-2803 and EIS 1075-V2.1
- DTE-initiated calls
- Host-initiated calls

Protocol summary

This section provides a brief summary of the T3POS protocol. For complete details on the protocol and the MAX X.25 PAD, refer to the documents listed in “References” on page 225.

The T3POS protocol provides reliable and efficient data interchange (transactions) between a host (usually a transaction server) and a DTE (usually a client). The T3POS DTE is usually a client device communicating through an asynchronous port, while the T3POS host is a mainframe or server communicating through an X.25 packet network. The T3POS PAD (the MAX) converts data arriving from a T3POS DTE to a format that is capable of being transmitted over a packet network. It also ensures reliability and efficiency as described in the protocol.

Note that the T3POS PAD does not alter, check or convert the parity of characters it receives from or sends to the X.25 network or the T3POS DTE. T3POS essentially uses a data format of 8 bit no parity, or more accurately 7 bits, 1 parity, but the parity bit is ignored

T3POS frame types

Depending on the current state of a transaction or call, and the mode of operation selected, T3POS uses different data formats and frame structures. The MAX supports four modes of operation: Local, Binary-local, Transparent and Blind.

General frames

A general frame (also known as a data frame) is defined as any sequence of octets received from or sent to the DTE within the period specified by the T1 timer (this timer is known as the char-to-char timer). Furthermore, in Local and Binary-local modes and in opening frames, general frames are encapsulated in the format:

<STX [data] ETX XRC>

where:

- STX is the ascii character \002
- data is the user data being sent in this frame
- ETX the ascii character \003
- XRC is the checksum.

For all modes except Binary Local the checksum is a one character Longitudinal Redundancy Check (LRC) checksum. For Binary Local mode, the checksum is a two character Cyclic Redundancy Check (CRC) checksum.

Control frames

Control frames are used only when a call is being established and not during data transfer. You can configure the T3POS modes and most of the T3POS parameters for the T3POS PAD using the VT-100 interface in the MAX. However, the operating mode as well as called number and call user data and some user facilities can be overridden by using a control frame. A control frame is a supervisory frame of the format:

<SOH MSS CUD STX [data] ETX XRC>

where:

- SOH is the ascii character \001
- MSS is the Mode Selection Signal that can be (optionally) used to indicate the mode for the call
- CUD is the Called User Data
This may contain an X.121 address, and user facilities or call user data in an X.28 format.
- data is optional in the control frame. In Transparent and Blind modes, the T3POS PAD is essentially restricted to passing data frames between the T3POS DTE and the T3POS host.
- ETX the ascii character \003
- XRC is the checksum.

For all modes except Binary Local the checksum is a one character Longitudinal Redundancy Check (LRC) checksum. For Binary Local mode, the checksum is a two character Cyclic Redundancy Check (CRC) checksum.

T3POS Timers

The T3POS protocol defines six timers:

- T1: Char-to-Char timer
- T2: SYN-to-SYN timer
- T3: ENQ Handling timer
- T4: Response timer
- T5: DLE, EOT timer
- T6: Frame Arrival timer

DTE-initiated calls

If the first T3POS frame (which can be either a general frame or a control frame) the MAX receives is from the DTE, the session is qualified as DTE-initiated. When the MAX receives a general frame from the DTE it triggers a call to the host using the settings in the Answer profile (or the Connection profile). When the MAX receives a control frame from the DTE it also triggers a call to the host. In this case, however, the mode and called address specified in the control frame (if any) is used for the call, overriding any setting configured in the MAX.

Host-initiated calls

This implementation does not directly support incoming calls to the DTE. Instead, calls initiated by the host are answered by the DTE connecting to the T3POS PAD and “listening” for host-initiated calls. The host must send a called address matching the pattern the DTE is listening for. This pattern does not need to be a complete X.121 address, but could be a sub-pattern (including wildcard characters). You configure the listening pattern using the Listen X.121 Addr parameter described on page -235.

Flow control

Flow control should not be an issue for the X25 T3POS implementation. This is because the T3POS protocol has an effective window size of one (that is, every frame must be acknowledged before another frame is sent) and because the MAX buffers all the frames before forwarding them to the DTE or the host. However, you should chose the T2, T3 and T4 timers carefully to account for the fact that the MAX buffers the data before forwarding it. Note that the current Ascend modem code does RTS/CTS flow control all the time and this cannot be disabled.

References

The T3POS protocols derived from several documents that have become de facto standards:

- GR-2803: Generic requirements for a Packet Assembler/Disassembler supporting T3POS, Bellcore GR-2803-CORE Issue 2, Dec. 1995.
This is the basic defining document.
- EIS 1075-V2.1: External Interface Specification for Data-Terminal-Equipment support of T3POS, Applied Digital Design, version 2.1, March 1994.
Specifies error recovery mechanisms between a T3POS DTE and a T3POS PAD on one side and a T3POS PAD and the T3POS host in the other side.

Refer to the MAX 4.6C and 5.0A addenda for information on the MAX X25/PAD.

Configuring a T3POS connection

You can configure a T3POS connection using either the Connection profile (for authenticated users) or an Answer profile (for unauthenticated users). Refer to “New parameters” on page 228 for descriptions of the new T3POS parameters.

Configuring a T3POS connection consists of these general steps:

- Create a Connection profile or an Answer profile for the user connecting to the T3POS.
- Create an X.25 profile that defines the X.25 connection the T3POS PAD uses.

Note: The settings in the Connection or Answer profile can be overridden by the settings sent in control frames.

This section explains how to create a Connection profile or an Answer profile for T3POS. Refer to the MAX 4.6C and 5.0A addenda for information on configuring an X.25 profile.

To configure a T3POS Connection profile:

- 1 From the main Edit menu select Ethernet>Connections>*any Connection profile*.
- 2 Set Active to Yes.
- 3 Set Encaps to X25/T3POS.
- 4 Open the Encaps Options submenu.
- 5 Set X.25 Prof to the name of the X.25 that is to be used for this T3POS connection.
The X.25 profile must exist and be active before you can save this Connection profile.
- 6 Specify the Recv PW used to authenticate the caller.
- 7 Specify the parameters used for the T3POS connection. Refer to the “New parameters” on page 228 for descriptions of the T3POS parameters.
- 8 Exit and save the Connection profile.

To configure a T3POS Answer profile:

- 1 From the main Edit menu select Ethernet>Answer>Encaps.
- 2 Set X25/PAD to Yes and X25/T3POS to Yes.
- 3 Exit the Encaps submenu.
- 4 Select T3POS Options.
- 5 Set X.25 Prof to the name of the X.25 that is to be used for this T3POS connection.
The X.25 profile must exist and be active before you can save the Answer profile.
- 6 Specify the parameters used for the T3POS connection. Refer to the “New parameters” on page 228 for descriptions of the T3POS parameters.
- 7 Exit and save the Answer profile.

Accessing the T3POS

User can access the T3POS in any of the following ways:

- Through a modem (for MAX units only)

- Via a TCP/IP client to the default TCP modem port 6150 (or to the TCP modem port configured on the Ascend unit)
- Via a TCP/IP client to port 23 (for Telnet access) or to 513 (for rlogin access)

Accessing the T3POS from a dial-in connection

This following example describes how to access the X.25/T3POS from a modem. Note that the X.25 data link is already up because it is a nailed physical connection. This scenario also applies to Telnet users connecting the port 150 of the MAX.

Note that Telnet client programs should use 8 bit mode to connect to the MAX.

- 1 A user dials in through a modem or through Telnet.
- 2 The user is authenticated against a Connection profile. If no Connection profile exists for the user, the Answer profile is used (if configured).
- 3 Both the Connection and Answer profiles specify that the user is an X.25 user (that is, Encaps is set to X25/T3POS) as well as an X.25 profile that specifies the physical interface where the X.25 call is to be established.
The X.25 profile determines the settings for the LAPB (or LAPD) and packet level, including timers, window size, and so on. For LAPB, the X.25 profile also specifies the nailed group to use for the logical call.
- 4 The connection is then established using the settings in both the Connection profile (or Answer profile) and the X.25 profile and the call is directed to the T3POS.
- 5 The user then must use the normal X.25/PAD commands as explained in the MAX 4.6C and 5.0A addenda.

Accessing the T3POS from the MAX terminal server interface

This following example describes how to access the X.25/T3POS from the MAX terminal server interface or through Telnet.

- 1 From the terminal server prompt, the user enters the T3POS command. For example:
`ascend% t3pos`
- 2 The user is then directed to the T3POS PAD and T3POS traffic can now be transmitted.

Accessing the T3POS through immediate mode

To allow access the T3POS PAD immediately upon connecting, set Immediate Service to X25/T3POS in the Ethernet>Mod Config>TServ options submenu. This is typically how users will connect to the T3POS PAD.

We recommend that when use immediate service, you suppress the terminal server banner (using the Banner parameter) as well as reducing the PPP delay parameter to its minimum. Both these parameters are in the Ethernet>Mod Config>TServ options submenu.

User interface changes

This section describes the changes to the user interface for T3POS. Note that some of these parameters, such as the data transfer mode and the X.121 address, can be overridden by the opening control frame from the DTE according to the T3POS protocol.

Parameters not applicable for T3POS

The following parameters are not applicable when Encaps is set to T3POS

- IP routing

Changed parameters

In the Ethernet>Mod Config>Tserv options submenu, the Immediate Service parameter has been changed. An options for T3POS has been added. This is to support immediate T3POS mode.

New parameters

The following parameters appear in a Connection profile Encaps options submenu when you set Encaps to X25/T3POS as well as in the Answer profile X25/T3POS submenu.

```
s0-lnn name
Encaps options...
  X.25 Prof=
  RECV Password= (Connection profile only)
>Host init. mode=Local
  DTE init. mode=Local
  ENQ handling=Off
  Max. block size=512
  T3POS T1=5
  T3POS T2=40
  T3POS T3=15
  T3POS T4=40
  T3POS T5=2400
  T3POS T6=300
  Direct Call Addr X.121 addr=
  Method of host notif=None
  PID selection=X.29
  ACK suppression=Off
  Data Format=7-E-1
  Link Access Type=Dedicated
  Retry Limit=3
  Listen X.121 Addr=
  Reverse Charge=No
  RPOA=
  CUG Index=
  NUI=
```

X.25 Prof

Description: This is the name of an X.25 profile that carries the X.25 logical connections. If the matching X.25 profile can not be found, no session is started for this Connection profile. To guard against this misconfiguration, an active connection profile specifying X.25 encapsulation can not be saved unless the named X.25 profile has been defined and is active.

Usage: Enter the name of the X.25 profile used for this Connection profile. This can be up to 15 characters. The default is null.

Dependencies: This parameter is always applicable.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>T3POS options

Recv PW

Description: Specifies the password for the Connection profile.

Usage: Enter the name of the password this Connection profile. This can be up to 20 characters. The default is null.

Dependencies: This parameter is always applicable.

Location: Ethernet>Connections>*any Connection profile*>Encaps options

Host init. mode

Description: For host-initiated calls, this specifies the default data transfer mode. Note that the host can override this setting with a control frame.

Usage: Specify one of the following values:

- Local (the default)
Specifies that error recovery is performed locally. In this mode, the MAX does not send supervisory frames that is, ACKs and NAKs) across the X.25 network. The T3POS PAD is responsible for sending supervisor frames to the T3POS DTE.
- Transparent
Specifies that the T3POS PAD does not provide any error recovery. In this mode, the DTE and the host system provide error recovery for the connection. In Transparent mode, however, the T3POS PAD does recognize a clear request command signal from the DTE (that is, DLE, EOT) and clears the call when it receives a DLE, EOT command.
- Blind
The same as Transparent mode except that the T3POS PAD does not clear a call when it receives a clear request command from the DTE. In this mode, the PAD or the host system must clear the call. The PAD passes all data "blindly," without regard to the protocol in use. This mode provides a means to pass raw binary data between the DTE and the host system without reference to the protocol being used.
- Bin-Local
Specifies that there is no error recovery between the T3POS PAD and the host but that there is error recovery between the PAD and the DTE. Like Blind mode, it passes data

between the DTE and the host without reference to the protocol being used., but continues to use the T3POS protocol between the DTE and the PAD.

Dependencies: This parameter is always applicable.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>T3POS options

DTE init. mode

Description: For DTE-initiated calls, this specifies the default data transfer mode. Note that the DTE can override this setting with a opening frame.

Usage: Specify one of the following values:

- Local (the default)
Specifies that error recovery is performed locally. In this mode, the MAX does not send supervisory frames that is, ACKs and NAKs) across the X.25 network. The T3POS PAD is responsible for sending supervisor frames to the T3POS DTE.
- Transparent
Specifies that the T3POS PAD does not provide any error recovery. In this mode, the DTE and the host system provide error recovery for the connection. In Transparent mode, however, the T3POS PAD does recognize a clear request command signal from the DTE (that is, DLE, EOT) and clears the call when it receives a DLE, EOT command.
- Blind
The same as Transparent mode except that the T3POS PAD does not clear a call when it receives a clear request command from the DTE. In this mode, the PAD or the host system must clear the call. The PAD passes all data "blindly," without regard to the protocol in use. This mode provides a means to pass raw binary data between the DTE and the host system without reference to the protocol being used.
- Bin-Local
Specifies that there is no error recovery between the T3POS PAD and the host but that there is error recovery between the PAD and the DTE. Like Blind mode, it passes data between the DTE and the host without reference to the protocol being used., but continues to use the T3POS protocol between the DTE and the PAD.

Dependencies: This parameter is always applicable.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>T3POS options

ENQ handling

Description: Specifies whether the PAD should expect to receive an ENQ from the host when an X.25 virtual call is established. ENQ indicates that the host is ready to receive data.

Usage: Specify one of the following values:

- Off (the default)
Specifies that the PAD does not expect to receive an ENQ before sending data to the host. The host is ready to receive data as soon as the X25 call is established.
- On

Specifies that the PAD does not forward data it receives from the DTE to the host until it either receives an ENQ or the T3 POS timer expires. Note that the PAD does not forward the ENQ to the DTE.

Dependencies: This parameter is always applicable.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>T3POS options

Max. Block Size

Description: Specifies maximum length of a transmission (including the length of opening frame) in bytes that the PAD must be able to accept and process from the DTE or host. This only applies to processing opening frame and to both local modes of operation.

Usage: Specify one of the following values:

- 512 (the default)
- 1024

Dependencies: Keep this additional information in mind.

- The Max. Block Size may apply even if both the host and DTE initiated call default mode are non-local. This is because the mode can be changed through an opening frame, in which case this parameter applies.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>T3POS options

T3POS T1

Description: Specifies the Char-to-Char timer. This timer indicates the maximum amount of time permitted between characters sent from the DTE to the PAD.

Usage: Specify a value between 1 and 20 (tenths of seconds). The default is 5.

Dependencies: This parameter is always applicable.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>T3POS options

T3POS T2

Description: Specifies the SYN-to-SYN timer. This timer applies to opening frames in Local or Bin-Local mode. Normally, the PAD sends SYN signals to the DTE at the interval specified by the T2 timer to indicate that an idle link is still alive. However, if the DTE sends a SYN signal to the PAD before the PAD sends one to the DTE, the T2 timer specifies the period of time the PAD expects SYN signals from the DTE. If the PAD does not receive two SYN signals with the interval specified by the T2 timer, it tries to restore the link.

Usage: Specify a value between 10 and 100 (tenths of seconds). The default is 40.

Dependencies: Keep this additional information in mind.

- The T2 timer only applies to the opening frame and to Local or Bin-Local mode.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>T3POS options

T3POS T3

Description: Specifies the ENQ handling timer. This timer indicates the amount of time the PAD waits for an ENQ from the host.

Usage: Specify a value between 5 and 50 (tenths of seconds). The default is 15.

Dependencies: Keep this additional information in mind.

- This is not applicable when you set ENQ Handling to Off.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>T3POS options

T3POS T4

Description: Specifies the Response Timer. This timer indicates the amount of time the PAD waits for a SYN from the DTE while the PAS is waiting for a response from the DTE. The SYN signal indicates that the response from the DTE is being delayed and also indicates that the link is still alive.

Usage: Specify a value between 10 and 100 (tenths of seconds). The default is 40.

Dependencies: This parameter is always applicable.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>T3POS options

T3POS T5

Description: Specifies the DLE, EOT timer. This timer indicates the maximum idle-time the PAD allows for a T3POS call (this is similar to the VC inactivity timer in the X25/PAD). The T5 timer applies only to transparent and blind mode only; it is disabled in both Local mode and Bin-Local mode.

Usage: Specify a value between 50 and 3000 (tenths of seconds). The default is 2400 (four minutes).

Dependencies: Keep this additional information in mind.

- The T5 timer may apply even if the default modes for both the host- and DTE-initiated calls are Local or Bin-Local. This is because the mode can be changed through an opening frame, in which case this parameter applies.
- The T5 timer applies only to transparent and blind mode only; it is disabled in both Local mode and Bin-Local mode.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>T3POS options

T3POS T6

Description: Specifies the Frame Arrival timeout. This timer indicates the maximum amount of time allowed between the time a dial-up connection is established and the first character of an opening frame is received.

Usage: Specify a value between 50 and 3000 (tenths of seconds). The default is 300 (30 seconds).

Dependencies: This parameter is always applicable.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>T3POS options

Direct Call Addr

Description: For DTE-initiated calls, this specifies the default host's X.121 address.

Usage: Specify an alphanumeric string. You can enter up to 15 characters. The default is null.

Dependencies: This parameter is always applicable.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>T3POS options

Method of host notif

Description: For DTE-initiated calls, this specifies how the host is notified of the mode of the call.

Usage: Specify one of the following values:

- None (the default)
Specifies that the host is not notified of the mode of the call and any data in the CUD is discarded.
- CRP
Specifies that the host is informed of the mode of the call by the DTE sending a Call Request Packet (CRP) in the CUD field of a control frame.
- MSF
Specifies that the host is informed of the mode of the call by the DTE sending a Mode Switch Frame (MSF) after the call has been established.

Dependencies: Keep this additional information in mind.

- This parameter does not apply if the opening frame is a general frame. In this case the default DTE-initiated mode is not changed.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>T3POS options

PID selection

Description: For DTE-initiated calls, this specifies which Protocol Identifier (PID) the PAD includes in the call request packet it sends to the host.

Usage: Specify one of the following values:

- X.29 (the default)
Specifies that the PAD sets the protocol identifier in the CUD field to X.29.
- T3POS
Specifies that the PAD sets the protocol identifier in the CUD field to T3POS.

Dependencies: This parameter is always applicable.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>T3POS options

ACK Suppression

Description: For DTE-initiated calls, this specifies whether the PAD sends an acknowledgment when it receives an opening frame from the DTE and also when it establishes a virtual call with the host.

Usage: Specify one of the following values:

- Off (the default)
Specifies that the PAD acknowledges the DTE's opening frame and the establishment of a call with the host.
- On
Specifies that the PAD does not acknowledge either the DTE's opening frame or the establishment of a call with the host.

Dependencies: Keep this additional information in mind.

- ACK Suppression only applies to DTE-initiated calls using Transparent or Blind mode.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>T3POS options

Data Format

Description: Specifies the data format and parity checking/generation behavior of the PAD when it validates opening frames as well as during Local mode data transfer.

Usage: Specify one of the following values:

- 7-E-1 (the default)
Specifies that the PAD uses 7 data bits, even parity, and 1 stop bit during opening frame validation and local mode data transfer.
- 7-O-1
Specifies that the PAD uses 7 data bits, odd parity, and 1 stop bit during opening frame validation and local mode data transfer.
- 7-M-1
Specifies that the PAD uses 7 data bits, mark parity, and 1 stop bit during opening frame validation and local mode data transfer.
- 7-S-1

Specifies that the PAD uses 7 data bits, space parity, and 1 stop bit during opening frame validation and local mode data transfer.

- 8-N-1

Specifies that the PAD uses 8 data bits, no parity, and 1 stop bit during opening frame validation and local mode data transfer.

Dependencies: This parameter is always applicable.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>T3POS options

Link Access Type

Description: Specifies the type of the DTE connection.

Usage: Specify one of the following values:

- Dedicated (the default)
Specifies that the DTE connection is a permanent, leased-line connection.
- Dial
Specifies that the DTE connection is a dial-up connection.

Dependencies: This parameter is always applicable.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>T3POS options

Retry limit

Description: Specifies the number of times in a row, per connection, that the PAD allows the DTE to send a frame or frame acknowledgment in error before it disconnects the call. For a dial-up connection, the Retry Limit specifies how many times the PAD will allow the DTE to try to establish a call that fails because the X.25 virtual call to the host could not be established. When the DTE exceeds the Retry Limit, the PAD disconnects the call.

Usage: Specify a value between 1 and 15. The default is 3.

Dependencies: This parameter is always applicable.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>T3POS options

Listen X.121 Addr

Description: Specifies a listen pattern for host-initiated calls. This is similar to typing the following command in the X.25 PAD:

```
* listen addr=pattern
```

The pattern is in the same format as an X.121 address, or sub address and can contain wild cards.

Usage: Specify an address. You can enter up to 15 characters.

Dependencies: This parameter is always applicable.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>T3POS options

Reverse Charge

Description: Specifies whether the call packet should include a reverse charge request facility parameter.

Usage: Specify one of the following values:

- Yes
Specifies that the call packet includes a reverse charge request facility parameter.
- No (the default)
Specifies that the call packet does not include a reverse charge request facility parameter.

Dependencies: This parameter is always applicable.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>T3POS options

RPOA

Description: Specifies the set of Recognized Private Operating Agency (RPOA) user facilities to use in the next call request. The RPOA facilities provide the data network identification code for the requested initial RPOA transit network and is in the form of four decimal digits.

Usage: Specify the RPOA user facilities to use in the next call request. You can specify up to four digits. The default is null.

Dependencies: Encaps must be set to X25/PAD for RPOA to be applicable.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>PAD options
Ethernet>Answer>T3POS options

CUG Index

Description: Specifies the closed user group (CUG) index/selection facility to use in the next call request. The closed user group selection/index facility is used to indicate to the called switch the closed user group selected for a virtual call.

Usage: Specify the CUG Index to use in the next call request. You can specify up to two digits. The default is null.

Dependencies: Encaps must be set to X25/PAD for CUG Index to be applicable.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>PAD options
Ethernet>Answer>T3POS options

NUI

Description: Specifies the set of Network User Identification (NUI) related facilities to use in next call request. NUI provides information to the network for purpose of billing, security, network management, or to invoke subscribed facilities.

Usage: Specify the NUI to use in the next call request. You can specify up to six digits. The default is null.

Dependencies: Encaps must be set to X25/PAD for NUI to be applicable.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>PAD options
Ethernet>Answer>T3POS options

X.25 over the D channel

This feature provides X.25 support over the ISDN D channel.

In this release, you can run existing X.25 applications, X.25 PAD, IP over X.25, and X.25 T3POS, over the D channel. For information on configuring X.25 on the MAX, refer to the 4.6C and 5.0A Release Notes or Addenda.

New parameters

This section describes the new parameters.

TEI

Description: Specifies the Terminal Endpoint Identifier (TEI). Your service provider can provide you with the appropriate value.

Usage: Specify a TEI value from 0 to 63. The default value is 23. If you set TEI to 0, the Ascend unit requests a TEI assignment from the network.

Location: Ethernet > X.25 > *any X.25 profile*

Changed parameters

This section describes the parameters that have changed to support X.25 over the D channel.

Call Type

Description: A new value, D-Channel, is now available. Select D-Channel if you want to run X.25 applications over the D channel.

Dependencies: The following parameters in the X.25 profile are not applicable when you set Call Type to D-Channel:

- Nailed Grp
- Data Svc

X.25 features

X.29 Reselection PAD message support

- PRI # Type
- Dial #
- Bill #
- Call-by-Call
- Transit #
- LAPB T1
- LAPB T2
- LAPB N2
- LAPB K
- X.25 Seq Number Mode
- X.25 Link Setup Mode
- X.25 Node Type
- X.25 Pkt Size
- X.25 Min Pkt Size
- X.25 Max Pkt Size

X.29 Reselection PAD message support

The MAX now supports the X.29 reselection message. This can be used to allow MAX to accept all incoming point-of-sale (POS) transactions and then redirect them to the appropriate authorization host. This is transparent to the user connecting to the MAX.

Overview

This feature adds support for X.29 reselection messages. This feature is typically used for credit card authorizations. A POS terminal connects to the MAX, which then establishes a connection to a central host. The central host determines to which authorization host the call should be directed. It then sends an X.29 reselection message back to the MAX with the X.25 Data Network Address (DNA) of the authorization host. When the MAX receives this message, it drops the connection to the central host and establishes a connection to the correct authorization host. This is transparent to the user at the POS terminal.

Answer Profile for X25/PAD terminal server users

X25/PAD users operating the PAD via the terminal server prompt or the terminal server's immediate X25/PAD services are not authenticated; therefore, they have no Connection profile associated with them. In past releases, such users would have hard-coded X.25 defaults assigned to them. In this release, you can configure X.25 options in the Answer profile to provide customized defaults.

User interface changes

The X25 Options submenu now appear below the Answer menu. It contains these parameters:

- LCN
- X.3 Param Prof
- Max Unsucc. calls
- VC Timer enable
- Auto-Call X.121 addr
- Reverse Charge
- X.3 Custom

Each of these parameters has the same functionality as the Connection profile parameter of the same name. For information on all other options, see the previous release notes.

Specifying X.3 parameters in an X.25 PAD Connection profile

Ascend's X.25/PAD implementation contains ten permanent X.3 parameter profiles that contain settings for a range of devices, such as terminals and printers. However, these profiles do not cover all devices. In this release, you can define and save a new profile that describes a device not specified by the permanent profiles.

User interface changes

A new parameter, X.3 Custom, appears in both the Answer profile and the Connection profile. The following section describes this new parameter.

X.3 Custom

Description: This parameter specifies a string containing X.3 profile parameters. The Ascend unit parses this string into X.3 profile parameters when an operator uses the PAD.

Usage: Specify a string using this format:

X.3 Custom=[<ref>:]<val>,[<ref>:]<val>,...,[<ref>:]<val>

The <ref> argument is the number of an X.3 parameter as defined in the ITU X.3 specification. You can specify a value between 1 and 22. By default, the <ref> value starts at 1 and is incremented by 1 after each comma. Unless you wish to specify fewer X.3 parameters than the maximum, you need not enter the <ref> argument.

The <val> argument is the value associated with the X.3 parameter.

The Ascend unit silently ignores invalid parameters.

You can enter up to 64 characters for the entire X.3 Custom specification. By default, the X.3 Custom parameter contains the X.3 parameter values set in the CRT profile.

Dependencies: The X.3 Custom parameter does not apply if X.3 Param Prof is not set to CUSTOM.

X.25 features

X.25/IP inactivity timer supported

Location: Answer profile: Ethernet>Answer>X.25 Options

Connection profile: Ethernet>Connections>Any Connection profile>Encaps Options

See Also: X.3 Param Prof

X.25/IP inactivity timer supported

The MAX now supports an inactivity timer for IP over X.25.

The MAX now supports an inactivity timer for X.25/IP calls. This new parameter is described below.

Inactivity Timer

Description: The inactivity timer specifies the number of seconds to allow a connection to remain inactive before dropping the virtual circuit.

Usage: Specify a number of seconds. The default zero disables the inactivity timer.

Example: Inactivity Timer=120

Dependencies: This parameter applies only to X.25/IP connections

Location: Ethernet>Connections>Encaps Options

New X.25 user facilities parameters added

Three new user facilities parameters have been added to MAX units that support X.25.

This feature provides X.25 users with access to three more user facility parameters. Previously, users could only access these features using the `fac` command in the X.25 PAD. Now these parameters are now available to immediate X.25 PAD and T3POS users.

Parameter reference

RPOA

Description: Specifies the set of Recognized Private Operating Agency (RPOA) user facilities to use in the next call request. The RPOA facilities provide the data network identification code for the requested initial RPOA transit network.

Usage: Specify the RPOA user facilities to use in the next call request. You can specify up to four decimal digits. The default is null.

Dependencies: You must set Encaps to X25/PAD for RPOA to apply.

Location: Ethernet>Connections>*any Connection profile*>Encaps options

Ethernet>Answer>PAD options

Ethernet>Answer>T3POS options

CUG Index

Description: Specifies the closed user group (CUG) index/selection facility to use in the next call request. The closed user group selection/index facility specifies to the called switch the closed user group selected for a virtual call.

Usage: Specify the CUG Index to use in the next call request. You can specify up to two digits. The default is null.

Dependencies: You must set Encaps to X25/PAD for CUG Index to apply.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>PAD options
Ethernet>Answer>T3POS options

NUI

Description: Specifies the set of Network User Identification (NUI) related facilities to use in next call request. NUI provides information to the network for billing, security, network management purposes, and for activating subscribed facilities.

Usage: Specify the NUI to use in the next call request. You can specify up to six digits. The default is null.

Dependencies: You must set Encaps to X25/PAD for NUI to apply.

Location: Ethernet>Connections>*any Connection profile*>Encaps options
Ethernet>Answer>PAD options
Ethernet>Answer>T3POS options

Customized features

Personal Handy Phone Service (PHS) support (Japan only)

Personal Handy Phone Service (PHS) is a new feature that provides access for users of the PHS mobile phone system, currently available in Japan only. This new feature is available through the introduction of a new slot card and supporting software.

Overview

PHS is a mobile phone service currently offered in Japan only. In addition to voice communication, PHS offers data communication at bandwidth up to 32 Kilobits per second. You can use this service for phone calls as well as Internet access.

This new feature is available through the addition of a slot card, allowing 16 concurrent PHS users.

In addition, you need to enable the software functionality on the MAX, through a process commonly referred to as a Hash Code upgrade. Please see your Ascend Sales Representative with questions about the upgrade process.

When the MAX is booted with a PHS card in slot 4 and the software enabled, the following is displayed:

```
Main Edit Menu
  00-000 System
  10-000 Net/T1
  20-000 Net/T1
  30-000 Empty
  40-000 PIAFS-16
  50-000 Empty
  60-000 Empty
  70-000 Empty
  80-000 Empty
  90-000 Ethernet
  A0-000 Ether Data
  B0-000 Serial WAN
```

PIAFS stands for Personal Internet Access Forum Standard. PIAFS is a protocol designed to support connection negotiation, data transfers and error correction. The -16 refers to the slot card's support of 16 concurrent PHS users.

T-Online

RADIUS bootup server supported for ZGR subaddresses and answer numbers

You can now configure a special RADIUS bootup server to load the subaddresses and answer numbers the MAX needs for redirecting calls to a Deutsche Telekom ZGR.

Background

A ZGR is a piece of older equipment that Deutsch Telekom uses to give clients access to X.25 services. In a previous release, you could enable the MAX to act as a network switch, redirecting calls to a ZGR that handles T-Online services. (T-Online is Deutsche Telekom's online service.) To set up the MAX to redirect calls, you configured RADIUS to provide the MAX with the necessary ZGR subaddresses and answer numbers. An enhancement in the new release is a special RADIUS bootup server that loads both the subaddress list and the answer number list. The MAX does not use this special server for authentication.

Loading ZGR subaddresses

RADIUS loads ZGR subaddresses by means of a pseudo-user profile. For a unit-specific configuration, the first line has the following format:

```
DirDoSub-unit_name-num Password="Ascend", User-Service=Dialout-Framed-User
```

where *unit_name* is the system name of the MAX (the name specified by the Name parameter in the System profile) and *num* is a number in a sequential series starting at 1. For a global configuration, the first line has the following format:

```
DirDoSub-num Password="Ascend", User-Service=Dialout-Framed-User
```

After the first line, you specify one or more ZGR subaddresses by setting the Client-Port-DNIS attribute as many times as necessary.

For example, suppose that a MAX connects to a ZGR, and that you want the MAX to pass to the ZGR ports all calls that provide the subaddresses 1111, 1234, and 1982. For a MAX named Munich1, you would create the following pseudo-user profile:

```
DirDoSub-Munich1-1 Password="Ascend", User-Service=Dialout-Framed-User
      Client-Port-DNIS="1111",
      Client-Port-DNIS="1234",
      Client-Port-DNIS="1982"
```

Loading ZGR answer numbers

RADIUS loads ZGR answer numbers by means of a pseudo-user profile. For a unit-specific configuration, the first line has the following format:

```
DirDoNum-unit_name-num Password="Ascend", User-Service=Dialout-Framed-User
```

where *unit_name* is the system name of the MAX (the name specified by the Name parameter in the System profile) and *num* is a number in a sequential series starting at 1.

For a global configuration, the first line has the following format:

```
DirDoNum-num Password="Ascend", User-Service=Dialout-Framed-User
```

After the first line, you specify one or more ZGR answer numbers by setting the Client-Port-DNIS attribute as many times as necessary.

For example, suppose that a MAX connects to a ZGR, and that you want the MAX to pass to the ZGR ports all calls that provide the answer numbers 1111, 1234, and 1982. For a MAX named Munich1, you would create the following pseudo-user profile:

```
DirDoNum-Munich1-1 Password="Ascend", User-Service=Dialout-Framed-User
      Client-Port-DNIS="1111",
      Client-Port-DNIS="1234",
      Client-Port-DNIS="1982"
```

User interface changes

To support the RADIUS bootup server, the Ethernet > Mod Config > Auth menu now contains three new parameters. This release also contains new Dirdo terminal server commands for managing the answer number and subaddress lists. The commands enable you to show, add, or

delete entries. Finally, this release includes a trap that alerts you if the MAX does not receive its answer numbers and subaddresses.

New parameters

This release adds the parameters Auth Boot Host #1, Auth Boot Host #2, and Auth Boot Port to the Ethernet > Mod Config > Auth menu.

Auth Boot Host #1

Description: Specifies the IP address of the first RADIUS bootup server the MAX contacts, at startup, to obtain ZGR subaddresses or answer numbers.

Usage: Specify an IP address in dotted decimal notation. The default value is 0.0.0.0. If you accept the default, the MAX does not use a RADIUS server for ZGR subaddresses or answer numbers.

Dependencies: You can use the ZGR subaddress and answer number feature without specifying a special bootup server. If you don't specify a special bootup server, the MAX uses the authentication server specified by Auth Host in the Ethernet > Mod Config menu to store the ZGR subaddresses and answer numbers.

If you set the Auth Boot Host #1 parameter, you must also specify a value for the Auth Key and Auth Src Port parameters in the Ethernet > Mod Config > Auth menu.

Location: Ethernet > Mod Config > Auth

See Also: Auth Boot Host #2, Auth Boot Port

Auth Boot Host #2

Description: Specifies the IP address of the RADIUS server the MAX contacts if the server specified by Auth Boot Host #1 fails to respond.

Usage: Specify an IP address in dotted decimal notation. The default value is 0.0.0.0. If you accept the default, the MAX does not use a secondary RADIUS server for ZGR subaddresses or answer numbers.

Dependencies: You can use the ZGR subaddress and answer number feature without specifying a special bootup server. If you don't specify a special bootup server, the MAX uses the authentication server specified by Auth Host in the Ethernet > Mod Config menu to store the ZGR subaddresses and answer numbers.

If you set the Auth Boot Host #2 parameter, you must also specify a value for the Auth Key and Auth Src Port parameters in the Ethernet > Mod Config > Auth menu.

Location: Ethernet > Mod Config > Auth

See Also: Auth Boot Host #1, Auth Boot Port

Auth Boot Port

Description: Specifies the port number to use when contacting the RADIUS server specified by Auth Boot Host #1 or Auth Boot Host #2.

Usage: Specify a value between 0 and 1024. The default value is 0 (zero), which disables the RADIUS bootup-server feature.

Location: Ethernet > Mod Config > Auth

See Also: Auth Boot Host #1, Auth Boot Host #2

New terminal server commands

This release supports the new Dirdo commands, which enable you to show, add, or delete entries from the answer list or the subaddress list. Table 32 lists the new commands. To use them, you must have administrative authorization.

Table 32. Dirdo commands

Command	Description
Dirdo show ans sub	Lists all the answer numbers (when you specify ans) or all the subaddresses (when you specify sub) on the RADIUS bootup server.
Dirdo add ans <i>num</i> sub <i>num</i>	Adds the answer number (when you specify ans) or subaddress (when you specify sub) that you enter as the <i>num</i> argument. For example, to add the subaddress 1234 to the list, enter the following command: Dirdo add sub 1234
Dirdo del ans <i>num</i> sub <i>num</i>	Deletes the answer number (when you specify ans) or subaddress (when you specify sub) that you enter as the <i>num</i> argument. For example, to delete the subaddress 1234 from the list, enter the following command: Dirdo del sub 1234

New trap

This release includes a new trap with the value of TRAP_DIRDO_FAILURE (21). This trap is triggered when the MAX receives an incoming call without a matching answer number or subaddress. It has meaning only if you are using T-Online functionality.

ZGR answer numbers obtained from RADIUS

This release enables the MAX to act as a network switch, redirecting calls to a ZGR that handles online services. A ZGR is a piece of older equipment that Deutsch Telekom uses to give clients access to X.25 services. In this release, you can configure RADIUS to provide the MAX with the necessary ZGR answer numbers.

Configuring RADIUS with ZGR answer numbers

The MAX uses an answer number to determine whether it should forward a call to a ZGR. If the incoming number matches the answer number, the MAX forwards the call. To specify the ZGR answer numbers in RADIUS, follow these steps:

- 1 Create the first line of a pseudo-user profile using the User-Name, Password, and User-Service attributes.

You create a pseudo-user profile to store information that the MAX can query—in this case, to store ZGR answer numbers. You can configure pseudo-users for both global and MAX-specific configuration control of ZGR answer numbers. The MAX adds the unit-specific numbers in addition to the global numbers.

For a unit-specific configuration, specify the first line of a pseudo-user profile in this format:

```
DirdoNum-unit_name-num Password="Ascend", User-Service=Dialout-Framed-User
```

For a global configuration, specify the first line of a pseudo-user profile in this format:

```
DirdoNum-num Password="Ascend", User-Service=Dialout-Framed-User
```

unit_name is the system name of the MAX—that is, the name specified by the name parameter in the System profile. *num* is a number in a sequential series, starting at 1.

- 2 For each pseudo-user profile, specify one or more ZGR answer numbers using the Client-Port-DNIS attribute.

The combined total of all the answer numbers you define in pseudo-user profiles cannot exceed 100. This configuration constitutes a new use of Client-Port-DNIS, an attribute generally used for called number authentication.

For example, suppose that a MAX connects to a ZGR, and that you want the MAX to pass all calls that provide the answer numbers 1111, 1234, and 1982 to the ZGR ports. The pseudo-user profile for a MAX named Munich1 looks like this one:

```
DirdoNum-Munich1-1 Password="Ascend", User-Service=Dialout-Framed-User  
Client-Port-DNIS="1111",  
Client-Port-DNIS="1234",  
Client-Port-DNIS="1982"
```

How the MAX obtains ZGR answer numbers

Whenever you power on the MAX, reset the unit, update the RADIUS configuration with the Upd Rem Cfg command, or update RADIUS from SNMP, the MAX queries RADIUS for the ZGR answer numbers. RADIUS provides the answer numbers in this way:

- 1 RADIUS looks for profiles having the format DirdoNum-unit_name-1, where *unit_name* is the system name.

- 2 If at least one profile exists, RADIUS loads all existing profiles with the format `DirdoNum-unit_name-num`.
The variable *num* is a number in a sequential series, starting with 1.
- 3 The MAX queries for `DirdoNum-unit_name-1`, then `DirdoNum-unit_name-2`, and so on, until it receives an authentication reject from RADIUS.
- 4 RADIUS loads the global configuration profiles having the form `DirdoNum-num`.
- 5 The MAX queries `DirdoNum-1`, then `DirdoNum-2`, and so on, until it receives an authentication reject from RADIUS.

The MAX caches the answer numbers locally, and uses them to check against the answer numbers of incoming calls. Note that ports 3 and 4 are reserved for outbound connections to the ZGR device.

ZGR subaddresses obtained from RADIUS

This release enables the MAX to act as a network switch, redirecting calls to a ZGR that handles current services. A ZGR is a piece of older equipment that Deutsch Telekom uses to give clients access to X.25 services. In this release, you can configure RADIUS to provide the MAX with the necessary ZGR subaddresses.

Configuring RADIUS with ZGR subaddresses

To specify the ZGR subaddresses in RADIUS, you create a pseudo-user profile to store information that the MAX can query—in this case, to store ZGR subaddresses. You can configure a pseudo-user either for global configuration control of ZGR subaddresses, or for local control of the subaddresses specific to a MAX unit. The MAX adds both the unit-specific configuration and the global subaddresses to its routing table.

To configure the pseudo-user profile, follow these steps:

- 1 Create the first line using the User-Name, Password, and User-Service attributes.
For a unit-specific configuration, specify the attributes in this format:
`DirdoSub-unit_name-num Password="Ascend", User-Service=Dialog-Framed-User`
where *unit_name* is the system name of the MAX (the name specified by the Name parameter in the System profile) and *num* is a number in a sequential series, starting at 1.
For a global configuration, specify the attributes in this format:
`DirdoSub-num Password="Ascend", User-Service=Dialog-Framed-User`
- 2 For each pseudo-user profile, specify one or more ZGR subaddresses by setting the Client-Port-DNIS attribute as many times as necessary.
The combined total of all the subaddresses you define in pseudo-user profiles must not exceed 100. This configuration constitutes a new use of Client-Port-DNIS, an attribute generally used for called-number authentication.

For example, suppose that a MAX connects to a ZGR, and that you want the MAX to pass all calls that provide the subaddresses 1111, 1234, and 1982 to the ZGR ports. The pseudo-user profile for a MAX named Munich1 would look like this one:

`DirdoSub-Munich1-1 Password="Ascend", User-Service=Dialog-Framed-User`

```
Client-Port-DNIS="1111",  
Client-Port-DNIS="1234",  
Client-Port-DNIS="1982"
```

How the MAX obtains ZGR subaddresses

Whenever you power on the MAX, reset the unit, update the RADIUS configuration with the Upd Rem Cfg command, or update RADIUS from SNMP, the MAX queries RADIUS for the ZGR subaddresses. RADIUS provides the subaddresses in the following way:

- 1 RADIUS looks for profiles having the format DirdoSub-*unit_name*-1, where *unit_name* is the system name.
- 2 If at least one such profile exists, RADIUS loads all existing profiles that have the format DirdoSub-*unit_name-num*. The variable *num* is a number in a sequential series starting with 1.
- 3 The MAX queries for DirdoSub-*unit_name*-1, then DirdoSub-*unit_name*-2, and so on, until it receives an authentication reject from RADIUS.
- 4 RADIUS loads the global configuration profiles that have the form DirdoSub-*num*.
- 5 The MAX queries DirdoSub-1, then DirdoSub-2, and so on, until it receives an authentication reject from RADIUS.

The MAX caches the subaddresses locally, and uses them to check against the subaddresses of incoming calls. Note that ports 3 and 4 are reserved for outbound connections to the ZGR device.

DTPT encapsulation for T-Online PPP sessions

This release adds a new call management encapsulation type, DTPT, to support Deutsch Telekom's need for multiple simultaneous PPP sessions between an Ascend router and a T-Online access host (a ZGR).

Overview

In this release, you can use the MAX as a gateway for PPP clients. These clients can call into the MAX, be authenticated in the conventional manner, and then send IP packets whose destination address is a T-Online ZGR. Though in many respects the ZGR functions like an IP router using PPP, it can only accept packets from a single source address on any given B channel. Therefore, the MAX must establish a separate call, on a separate B channel, for each source address sending packets through the gateway. The MAX can make several simultaneous calls to one destination—T-Online's ZGR—but the link differs from an MP or MP+ session in the following ways:

- Each call appears to the ZGR as an independent PPP call from the MAX.
- Each call carries packets from a single source address.

In this release, for the ZGR, you can define a user profile that specifies the new DTPT encapsulation type so that each B-channel connection to the ZGR carries only one user's traffic.

User interface changes

This release includes the following user interface changes:

- In the Connection profile, you can specify DTPT for the Encaps value. You must set Encaps=DTPT for each Connection Profile that will connect to a ZGR.
- When you set Encaps=DTPT, the Encaps Options submenu contains a subset of the parameters found there when Encaps=PPP:
90-103 T-Online
Encaps options...
>Send Auth=PAP
Send PW=*****
MRU=1524
Link Comp=Stac
VJ Comp=Yes
You set each parameter as you would for any other profile. The parameters function identically to those in the Encaps Options submenu for PPP, except that you can specify only PAP for Send Auth.
- When Encaps=DTPT, AnsOrig is set to N/A in the Telco Options submenu.
- When Encaps=DTPT, the Private and RIP parameters are set to N/A (and internally forced to not advertise the route) in the IP Options submenu.

All parameters in all other submenus apply as they would for an ordinary PPP call. For example, you can set the Idle parameter in the Connection profile's Session Options submenu to specify the number of seconds the MAX waits before clearing a call when a session is inactive.

How DTPT connections work

The end user's connection to the MAX, the MAX unit's connection to the ZGR, the IP routing scheme, and the algorithms for call disconnects are in many respects similar to their non-DTPT counterparts. Security and reliability considerations vary with the authentication process in use.

End user's connection to the MAX

The T-Online user's connection to the MAX shares many characteristics of other, more conventional connections:

- The user can use PPP, MP, or MP+ to connect to the MAX. Only the MAX unit's connection to the ZGR limits each address to one B channel.
- The user establishes a PPP session on the MAX by either of the following means:
 - A digital PPP call
 - An analog call to the terminal server, followed by a command to enter PPP mode.
- Once having thereby established a PPP, MP, or MP+ session with the MAX, the user can receive any conventional routing services. The MAX appears to the user as a conventional access router, regardless of whether the destination is a ZGR. In fact, the user can access the ZGR and other destinations in the same session.

Customized features

DTPT encapsulation for T-Online PPP sessions

- A user not directly logged into the MAX, but with access to the LAN by other means, can forward IP packets via the MAX if the routing table is configured to permit it.
- If a user sends the MAX an IP packet with a destination of the ZGR, and the MAX does not have a call active to the ZGR, the MAX places a call to the ZGR and routes the subscriber's packet over the connection.
- The MAX forwards any traffic destined for a user with an active PPP session, including packets arriving from the ZGR.
- The MAX forwards any traffic with a destination on the LAN, including packets arriving from the ZGR.

The T-Online user's connection to the MAX differs from a conventional connection in the following ways:

- The MAX places a new call to the ZGR for each user on the basis of the user's IP address. That is, the MAX places a call for each source IP address, and sends only traffic from that address via the B channel to the ZGR.
- If the MAX has an active call to the ZGR, and additional packets for the ZGR arrive from the IP address for which the MAX initiated the call, the MAX sends the packets via the existing call.

MAX unit's connection to the ZGR

A call from the MAX to the ZGR is defined by a Connection profile, which in most respects resembles, and functions as, a conventional PPP dial-out profile:

- Each call from the MAX to the ZGR appears to the ZGR as a PPP call.
- Using the DTPT encapsulation type, the MAX internally manages a set of calls to the ZGR as a *bundle*, which is a single interface analogous to an MP or MP+ bundle.
- Authentication of the MAX-to-ZGR call takes place as defined in the Connection profile. The MAX reports its system name to the ZGR, and submits the password you specify using the Send PW parameter in the Encaps Options submenu of the Connection profile. For the Send Auth parameter, you can only select PAP.
- You can set the MRU, Link Comp, and VJ Comp parameters just as you would for a PPP call.
- The MAX dedicates particular E1 lines for calls to the ZGR. You direct calls to the appropriate line by using a trunk group (for switched calls) or a nailed group (for nailed-up calls). You can also use this mechanism to direct ZGR traffic to a particular B channel on a given E1 line.

The connection between the MAX and the ZGR differs from a conventional PPP connection in the following ways:

- The Connection profile uses DTPT to tell the MAX to perform outgoing call management according to the requirements of the ZGR.
- In contrast to MP or MP+, the MAX selects a channel on the basis of the source IP address. The MAX maintains a table of active sessions for this purpose. No Dynamic Bandwidth Allocation takes place, nor are packets statistically multiplexed among channels.
- During PPP negotiation with the ZGR, the MAX transmits the IP address of the user on whose behalf the call is being placed. (In this respect alone, the information the MAX sends to the ZGR is different from the information it supplies in an ordinary PPP call.)

Call disconnects

For a ZGR user, call disconnects work in the conventional manner in the following respects:

- The call from the user to the MAX is in all respects an ordinary PPP call. Therefore, the MAX can disconnect it for any of the ordinary reasons, such as the caller hanging up, the Connection profile's Idle parameter value being exceeded, or a lower layer failure.
- A call from the MAX to the ZGR can be disconnected for any of the reasons applicable to an ordinary PPP connection, including termination by the ZGR.
- Calls placed in response to packets received over the LAN do not automatically terminate when the source of the packets stops sending them, unless you set a value for the Idle parameter in the Connection profile.
- If you set in the Idle parameter in the Connection profile, and the MAX terminates the connection because the Idle value has been exceeded, the arrival of a subsequent packet from the associated user causes the MAX to re-establish the call to the ZGR.

For a ZGR connection, call disconnects have the following special features:

- If the connection to the user terminates for any reason, the MAX terminates the associated call to the ZGR.
- Although in some respects the set of calls to the ZGR resembles an MP or MP+ bundle, disconnection of one call does not imply disconnection of the entire bundle. They are separate PPP calls.
- No disconnection takes place based on Dynamic Bandwidth Allocation.

IP routing

IP routing works in the conventional manner in two respects. First, because the ZGR has one IP address, the MAX unit's routing table contains a single route to the ZGR, regardless of how many calls are actually active. As is the case for MP or MP+, the MAX selects a specific B channel appropriate to the encapsulation and call management type.

Second, with the exception of packets destined for the ZGR, the MAX routes any packets it receives from the user in the usual manner.

IP routing for the ZGR also differs from conventional routing in two ways. First, because every ZGR shares the same IP address, the MAX unit's route to the ZGR is not advertised. The Private and RIP parameters in the IP Options menu are N/A, causing the routes to remain private.

Second, packets received over a given WAN interface and destined for the ZGR must all have the same source address. More specifically, if the user's equipment is a gateway with other IP addresses behind it, equipment using those IP addresses cannot send packets to the ZGR. A call to the ZGR via the MAX has the same one-IP-address restriction as does a call directly to the ZGR. This restriction does not apply to packets received over the LAN.

Security and reliability considerations

Network integrity depends on the authentication process by which users gain access to the MAX. The MAX can gain access to the ZGR automatically. Because each B channel is reserved for one user at a time, the ZGR connection is subject to denial-of-service attacks if the MAX places calls to the ZGR without proper authentication. Call authentication covers access

to the MAX from the WAN; the design of the rest of the network must ensure security on the LAN.

Because the number of source addresses on the LAN can vary, up to a number larger than the number of interfaces on a MAX, it is impossible to guarantee that there will be an outgoing B channel available for every attempted access to the ZGR. In the presence of LAN-initiated calls to the ZGR, access might fail for LAN users, WAN users, or both. If you use the Idle parameter in the Connection profile, some users might obtain a connection to the ZGR on an initial attempt, but not on a reconnect attempt that occurs after the value of the Idle parameter has been exceeded.

PRI-PRI switching for T-Online

The current release provides PRI-PRI switching for T-Online. This feature provides a network-side implementation of NET-5 to support switching calls from Deutsche Telekom's public network to a T-Online server.

Overview

If T-Online is enabled, the MAX can switch calls from the public network to a T-Online server based on a match defined by Deutsche Telekom. The match can be one of the following:

- The caller's phone number matches a value specified in RADIUS, and the caller's subaddress matches a subaddress specified in the same RADIUS profile.
- The caller's phone number matches a value specified in RADIUS, and the caller does not have an associated subaddress.
- The caller's subaddress matches a value specified in RADIUS, and the caller does not have an associated phone number.
- The caller does not specify a subaddress or phone number.

User interface changes

This release includes the following interface changes:

- Two new parameters appear in the System Profile: T-Online and T302 Timer.
- In the Line Profile, the 1st Line and 2nd Line parameters accept new values: T-Online-USER and T-Online-ZGR.
- The Line Status menu displays new information based on the new values for 1st Line and 2nd Line.

The sections that follow describe each of these changes.

New parameters

T-Online

Description: This parameter specifies whether the MAX performs T-Online routing.

Usage: You can specify either Yes or No.

- Yes specifies that the MAX performs T-Online routing.
- No specifies that the MAX does not perform T-Online routing.
The default value is No.

Dependencies: If T-Online=Yes, you can not use lines 3 and 4 on the MAX for any purpose other than PRI-PRI switching.

Location: System Profile: System>Sys Config

See Also: T302 Timer

T302 Timer

Description: This parameter specifies the duration of the ISDN Q.931 layer 3 SETUP_ACK timer.

When the MAX receives the layer 3 SETUP message, the SETUP message consists of many IEs (Information Elements), such as Bearer Capability IE, Channel Identifier IE, Caller Number IE, Called Number IE, Sending Complete IE, and so on. The MAX checks for the Sending Complete IE upon receiving the SETUP message from the switch. If the Sending Complete IE is not in the SETUP message, the MAX starts the T302 timer and waits for an INFO message from switch. If the INFO message consists of Sending Complete IE, MAX stops the T302 timer. If no Sending Complete IE appears, the MAX restarts the T302 timer.

Usage: You can specify a value between 100 and 30000 one-hundredths of a second (1 to 30 seconds). The default value is 1800 (18 seconds).

Dependencies: T302 Timer does not apply if T-Online=No.

Location: System Profile: System>Sys Config

See Also: T-Online

Changed parameters

1st Line

Description: This parameter specifies how the E1 PRI interface for line #1 operates.

Usage: In addition to the existing settings, you can specify these new values for 1st Line:

- T-Online-USER
This setting indicates that the line connects to the switch, allowing the user to dial in.
- T-Online-ZGR
This setting indicates that the line connects to the ZGR server.

Location: Line profile: Net/E1>Line Config>Any Line profile

2nd Line

Description: This parameter specifies how the E1 PRI interface for line #2 operates.

Usage: In addition to the existing settings, you can specify these new values for 2nd Line:

- **T-Online-USER**
This setting indicates that the line connects to the switch, allowing the user to dial in.
- **T-Online-ZGR**
This setting indicates that the line connects to the ZGR server.

Location: Line profile: Net/E1>Line Config>Any Line profile

Status menu changes

For E1 lines configured for T-Online, the Line Stat window replaces the LA link status with new values:

- **NT**
This value specifies that the line's link status is up and it is physically connected to the ZGR server (1st Line or 2nd line parameter is set to T-Online-ZGR).
- **TE**
This value specifies that the line's link status is up and it is physically connected to the switch, allowing the user to dial in (1st Line or 2nd Line parameter is set to T-Online-USER).

For example, the T-Online Line Status display for line #1 in slot #1 which is connected to a ZGR server, would look like the following:

```
10-100 1234567890
  L1/NT  -----
3456789012345678901
---S-----
```

Appletalk features

AppleTalk routing added

An Ascend MAX or Pipeline can now function as an AppleTalk internet router, providing routing functions for AppleTalk nodes (Macintosh workstations or Apple printers) that are connected to the Ascend unit over AppleTalk Remote Access (ARA), Ethernet, or a WAN. Previously, Ascend units used bridging to provide AppleTalk connectivity.

To perform AppleTalk routing with your MAX or Pipeline, you should be sure the software load or Pipeline AppleTalk software you download contains the AppleTalk routing feature.

The new feature is designed for routing LAN traffic over synchronous PPP, MP, MPP, and Frame Relay WAN links. This implementation does not provide support for single-station dial-in over asynchronous PPP links such as modems. Currently, AppleTalk Remote Access (ARA) is the only way to dial into a remote MAX or Pipeline from a computer running the Mac OS.



Caution: Due to memory constraints, the AppleTalk software load disables firewalls on the Pipeline 50, Pipeline 75, and Pipeline 130.

Protocols implemented

MAX and Pipeline Appletalk routing support the following protocols:

- Datagram Delivery Protocol (DDP)
- Routing Table Maintenance Protocol (RTMP)
- AppleTalk Echo (AEP)
- Zone Information Protocol (ZIP)
- Name Binding Protocol (NBP)
- ATCP AppleTalk Control Protocol (ATCP—for router-to-router applications)

(For configuration instructions, see “Configuring AppleTalk routing” on page 259.)

Understanding network ranges and AppleTalk zones

AppleTalk routers provide for the configuration of network numbers and ranges and AppleTalk zones. Network numbers are assigned to network segments, and must be unique within the internetwork. A network range is a range of network numbers specified in the port descriptor of the router port and transmitted through RTMP to the other nodes of the network. Each of the numbers within a network range can represent up to 253 devices.

A zone is a multicast address containing an arbitrary subset of the AppleTalk nodes in an internet. Each node belongs to only one zone, but a particular extended network can contain nodes belonging to any number of zones. Zones provide departmental or other groupings of network entities that a user can easily understand.

In the Ascend AppleTalk router, zone names are not case sensitive. However, some routers regard zone names as case sensitive, and you should be consistent in spelling zone names when you configure multiple connections or routers. Although AppleTalk permits the use of spaces in zone names, it does not consider an underscore to be the same as a space. Since some routers do equate the underscore and the space, or do not recognize a space as a valid character, it is advisable to use only the underscore in a network with routers other than Ascend routers.

Extended and nonextended AppleTalk networks

AppleTalk uses two types of subnetworks: extended and nonextended. Nonextended networks theoretically allow up to 254 nodes. A nonextended network has one network number (not a range) and one zone. Examples of nonextended networks are LocalTalk and ARA dial-up networks.

An extended network is a group of nonextended networks on the same physical data link. It contains a range of network numbers, with each network in the range supporting up to 253 devices. EtherTalk and TokenTalk are examples of extended networks.

Seed routers

At least one router on a network, called a seed router, must have the network-number range in its port description. Other routers on the network can have a network range of 0, which means that they acquire the network-number range from RTMP packets sent by the seed router. To prevent conflicts, all seed routers on the same network must have the same value for the start and end of the network-number range.

Figure 1 shows a network with three routers and three zones configured. Each zone has a range of network numbers.

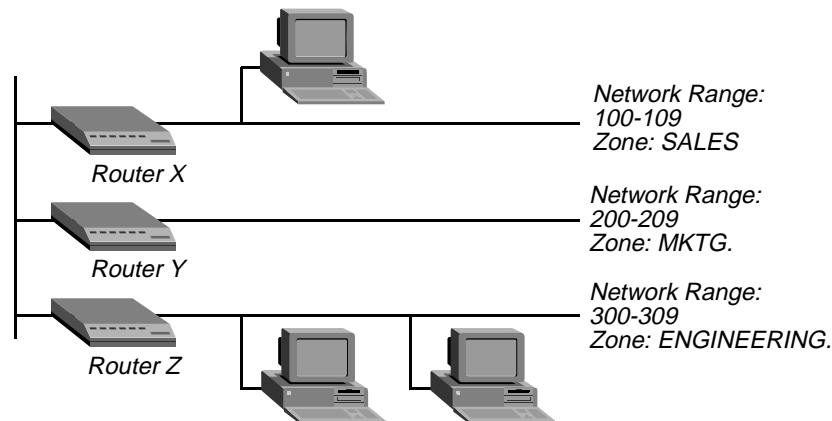


Figure 33. AppleTalk LAN

How AppleTalk works

Figure 2 shows a typical AppleTalk connection. The AppleTalk workstation is part of an Ethernet LAN connected to a Pipeline 75, which has a synchronous PPP WAN connection to a MAX 4000. One of the MAX 4000 ports is on a LAN that includes an Apple Laserwriter printer. Following is a brief, generalized description of how the workstation sends a file to the Laserwriter for printing:


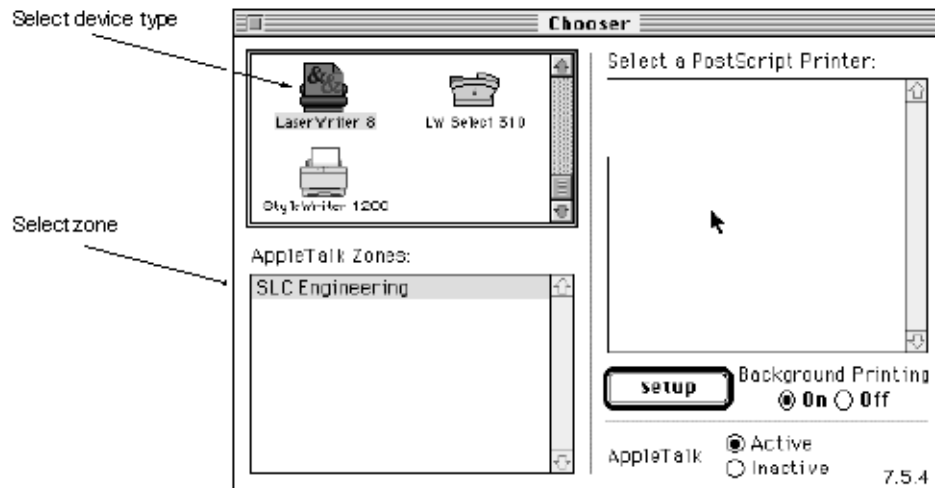


Figure 34. Routed connection

- 1 The AppleTalk workstation user opens the Macintosh Chooser.
The screen displays the network zones specified by the Connection profile stored in the Pipeline. (The first time a user opens the chooser, only the local Ethernet zones appear. That is, the WAN zone is the same as the local Ethernet zone.)
- 2 The Pipeline places the call and negotiates the WAN connection with the MAX 4000.
- 3 The workstation sends a ZIP Query to obtain an updated zone list from the MAX 4000, and the MAX returns the updated zone list. The new list, which might contain different zones from the initial zone list, replaces the initial list in the display and updates the Connection profile in the Pipeline.
- 4 The user selects a zone and a specific device in the Chooser.



- 5 The workstation sends a Name Binding Protocol (NBP) Broadcast Request to the Pipeline, which checks its Zone Information Table to identify the subnetwork in which that printer is located, and then sends the request to the MAX via the port configured in the Connection Profile.
- 6 The MAX determines the port to which the printer's subnetwork is attached, and looks up the printer by searching the multicast address assigned to the zone specified by the Pipeline.
- 7 All devices in the zone detect and process the NBP-lookup packet.
- 8 The selected printer obtains the sender's address from the lookup packet (sent by the workstation and forwarded by the routers), and sends the reply through the routers to the workstation.
- 9 The user sends the print job to the printer.
- 10 When the print job is complete and no data packets are passing through the connection, the MAX and the Pipeline continue to pass routing information until the idle timeout closes the connection. RTMP and ZIP packets do not reset the idle timer, but any other routeable packet to the network number or zone name specified for this connection does reset the timer.

After the link is dropped, the Pipeline retains in memory the last zone list displayed. If the workstation user opens the Chooser again, the list reappears and the process can begin again.

When to use AppleTalk routing

Use AppleTalk routing to connect two or more networks that have AppleTalk nodes such as Mac OS computers or Apple printers. Although you could use bridging, routing gives you more control over calls, reduces broadcast and multicast traffic over the WAN, and provides startup information to Appletalk nodes.

Call control

Bridging does not provide the call control that routing does, because bridging cannot detect NBP, RTMP, or ZIP packets, all of which can keep a connection open unnecessarily. If you use bridging with a filter to remove these packets from your WAN traffic, the routing information

becomes incorrect. A router does not have to bring up a line until it receives a data packet destined for the network number or zone name specified for a remote connection.

NBP packets

The Name Binding Protocol (NBP) drives call placement for the AppleTalk router. NBP is the protocol that enables AppleTalk users to issue a query (in the Chooser) for a type of service and receive specific server or printer names in response. For example, if a user selects LaserWriter in the Chooser and then selects a zone that is on a remote MAX, an NBP query goes to the router destined for the remote network that contains the selected zone. The router brings up a WAN connection in response to the NBP query, and places calls to connections that have AppleTalk routing enabled. The local AppleTalk router and the remote router exchange routing information until the idle timeout closes the connection.

RTMP packets

The AppleTalk router sends Routing Table Maintenance Protocol (RTMP) packets every 10 seconds. If you use bridging for AppleTalk, these packets keep a WAN link up for no other reason than to keep the RTMP information fresh across the WAN. If you add a filter to block RTMP packets, the zone and routing information age and become incorrect. But since RTMP information changes infrequently, it is not necessary to keep the WAN link up all the time. You really need only keep the information the same for the two ends of the WAN.

Once the connection is dropped (due to idle timeout), the Ascend AppleTalk router spoofs the RTMP and ZIP information on the Ethernet network until an event (such as a routeable packet to the network number or zone specified for this connection) requires bringing the WAN link up again. Until the event occurs, the user continues to the same zones see in the Chooser that were established by the connection, but the connection is not active.

RTMP (Routing Table Maintenance Protocol) and ZIP (Zone Information Packets) do not cause a call to be placed, and do not reset the idle timer. Any routeable packet to the network number or zone name specified for a connection brings the link back up.

Dropped connections in routed and bridged AppleTalk

If you filter RTMP packets in a bridged connection, you lose the AppleTalk network context when the connection closes at the end of a call. If another device comes onto the network while the first call is down, it might be assigned the same address as the device at the local end of the dropped call. Addressing inconsistency and network errors can result.

Because the Ascend AppleTalk routing implementation spoofs the RTMP and ZIP information on the Ethernet network, it can drop a call, freeing up the connection, without causing addressing inconsistencies.

Spoofing does not apply when an AppleTalk Filing Protocol volume is mounted on a workstation (selected in the Chooser). In this instance the server and workstation continue to exchange packets with numbers incremented serially, which would be very difficult to emulate or spoof.

Reducing broadcast and multicast traffic

Because AppleTalk uses multicast and broadcast addresses extensively, routing AppleTalk can greatly improve the efficiency of a LAN or WAN. By using AppleTalk zones to segment traffic, you can significantly reduce the amount of broadcast and multicast traffic on a LAN or WAN. When you set up a router for the first time, you identify the cable range (network number) for the subnetwork segment and one or more zones.

For example, when a user on a network without a router selects a device in the Chooser, the MAC OS computer sends out a Name Binding Protocol (NBP) Lookup as a broadcast packet. Since a bridge forwards all broadcast traffic, all devices on the network receive the Lookup packet. A router can significantly reduce AppleTalk traffic over the WAN because it does not forward broadcast traffic from one subnetwork to another, but stops it at the subnetwork port of the router.

Zone multicasting is intended to prevent any node not in the destination zone for the lookup from receiving the lookup packet. Any AppleTalk node responds only to NBP lookups for that node's zone name. In the example above, a router would convert the Broadcast Request packet generated by the Lookup request to a Forward Request packet for each network that contains nodes in the target zone specified by the Lookup request.

Routing vs. bridging multicast and broadcast traffic.

A bridge can filter directed traffic (traffic between two specific nodes) but cannot filter broadcast or multicast traffic, since there isn't a specific port that can be assigned to a multicast or broadcast address. This means that although filters used with bridging can reduce the number of AppleTalk packets sent to remote network segments, bridging does not reduce the number of broadcast and multicast packets over these networks.

Dynamic startup information

In addition to routing services, the Ascend AppleTalk router provides startup information to AppleTalk stations. Like other routed protocols, AppleTalk station or *node* addresses consist of a unique network-number/node combination. AppleTalk addresses are dynamically assigned when a node starts up. In addition, the router provides an AppleTalk node with the network-cable range to which it is attached and supplies zone-name information.

Configuring AppleTalk routing

To configure AppleTalk routing, you must set system-level parameters in the Ethernet Configuration profile and, if required for caller authentication, in the Answer profile. In addition, you can configure AppleTalk for specific connections.

Ethernet Configuration profile parameters

To set the required parameters in the Ethernet Configuration profile, open the Ether Options submenu of the Mod Config menu and:

- 1 Enable AppleTalk by setting the AppleTalk parameter to Yes. For example:

```
90-B00 Mod Config
  Ether Options...
    >Domain Name=abc.com
```

```
Pri DNS=200.00.200.14
Sec DNS=0.0.0.0/0
Pri DNS=0.0.0.0/0
Sec DNS=0.0.0.0/0
SNMP Options
Route Pref...
TServ options...
Bridging=No
IPX Routing=No
AppleTalk=Yes
Shared Prof=No
Telnet PW=
RIP Policy=Poisson Rvrs
RIP Summary=Yes
ICMP Redirect=Accept
DNS...
Auth
Accounting...
```

(You cannot perform the remaining steps of this procedure until you have set AppleTalk to Yes.)

- 2 In the AppleTalk Options submenu of the Ethernet Configuration profile, set the Zone Name parameter to the name of the zone to which the Ascend unit is connected. Enter up to 33 alphanumeric characters. For example, for router X in Figure 1:

```
90-B00 Mod Config
AppleTalk Options...
>Zone Name=SALES
AppleTalk Router=Seed
Net Start=300
Net End=309
Default Zone=SALES
Zone Name #1=MKTG
Zone Name #2=ENGINEERING
Zone Name #3=
Zone Name #4=
```

- 3 In the AppleTalk Options submenu of the Mod Config menu, set the AppleTalk Router parameter to Seed or Non-Seed to specify whether the Ascend unit is a seed or nonseed router. For example:

```
90-B00 Mod Config
AppleTalk Options...
>Zone Name=SALES
AppleTalk Router=Seed
Net Start=300
Net End=309
Default Zone=SALES
Zone Name #1=MKTG
Zone Name #2=ENGINEERING
Zone Name #3=
Zone Name #4=
```

A routed AppleTalk network must include at least one seed router. If you specify Non-Seed, the router learns network number and zone information from other routers. You can

set up more than one router on a network to be a seed router, but all seed routers must have the same value for both the start and end of the network number range.

The value zero (0) does not cause a conflict. Non-seed routers and other seed routers can have a value of 0 for the network number range. A router with a value of 0 for a network number range does not send this value to other routers, which means it does not seed the other routers in network with this range. The router with the zero value will not acquire a value for that network number range.

If you specify Non-Seed or Off, skip the remaining steps of this procedure.)

- 4 If the Ascend unit is a seed router, set the Net Start and Net End parameters to specify the range for the network to which the unit is attached. Valid values are from 1 to 65199. (For example, the menu shown in step 3 specifies a range of 300–309.)
(Remember that all seed routers on the same network must have the same network range.)
- 5 Specify the default-zone name for the nodes on the seed router's internet. Enter up to 33 alphanumeric characters in the Default Zone field. (For example, the menu shown in step 3 specifies SALES as the default zone.)
All AppleTalk nodes on the seeded network use the default zone until a user explicitly selects a different zone name.
- 6 Specify the names of the zones that the Ascend unit can seed. Enter up to 33 alphanumeric characters in each of one or more of the Zone Name fields. (For example, the menu shown in step 3 specifies MKTG in the Zone Name #1 field and SALES, MKTG in Zone Name #2.)

A Pipeline can seed up to five zones. A MAX can seed up to 32.

Answer profile parameter

If authentication uses names and passwords, enable AppleTalk routing in the Answer profile by selecting Route AppleTalk=Yes. For example:

```
90-700 Answer
PPP Options...
>Route IP=No
Route IPX=No
Route AppleTalk=Yes
Bridge=Yes
Recv Auth=None
MRU=1524
```

(You cannot set the Route AppleTalk parameter if AppleTalk is set to No in the Ethernet Configuration profile or if AppleTalk Router is set to Off in that profile's AppleTalk Options submenu.)

Configuring AppleTalk for a specific connection

To enable AppleTalk routing for a particular connection, open the Connection profile and:

- 1 Select Route AppleTalk=Yes. For example:

```
90-101 Macintosh 1
>Station=Macintosh 1
Active=Yes
Encaps=PPP
Dial # =1-100-111-1111
```

```
Route IP=No
Route IPX=No
Route AppleTalk=Yes
Bridge=No
Dial brdcast=N/A
Encaps options...
IP options...
IPX options
AppleTalk options...
```

(You cannot set the Route AppleTalk parameter if AppleTalk is set to No in the Ethernet Configuration profile or Route AppleTalk is set to No in the Answer profile.)

- 2 Set the Encaps option to PPP, MPP, or MP. (For example, the Connection profile shown in step 1 specifies PPP for the Encaps option.)
- 3 In the Dial # field, enter the number to call when this router encounters data destined for the remote Ascend AppleTalk router that is in the zone and network specified in the AppleTalk Options submenu.
- 4 Select AppleTalk Options to display the AppleTalk Options submenu
- 5 In the submenu's Zone Name field, enter up to 33 alphanumeric characters to specify the zone name for the AppleTalk router at the other end of the connection. For example:

```
90-101 Macintosh 1
>AppleTalk options...
Zone Name=ENGINEERING
Net Start=2001
Net End=2010
```

This zone name will appear in the AppleTalk Zones window of the Chooser. If the WAN segment for the zone is not already connected when packets for the zone are received (for example, when a user selects this zone in the Chooser, and then selects AppleShare), the Ascend unit places a call to the number in the Dial # field of the Connection Profile.

- 6 Enter the network range in the Net Start and Net End fields.
This range defines the networks available for packets that are to be routed to this static route. Valid entries for these fields are in the range from 1 to 65199. If there are other AppleTalk routers on the network, it is necessary to configure the network ranges to coincide with the other routers on the LAN.

Additional information about AppleTalk

This feature note provides only a very brief description of AppleTalk networking. For more complete information, see the following:

Apple Computer. *Inside Macintosh: Networking*.

Chappell, Laura A., and Roger L. Spicer. *Novell's Guide to Multiprotocol Internetworking*.

Sidhu, Andrews, and Alan B. Oppenheimer. *Inside AppleTalk, Second Edition*.

Cougias, Dell, and Heiberger. *Designing AppleTalk Network Architectures*.

Defender authentication added for AppleTalk Remote Access Protocol (ARAP)

A parameter, ARA, has been added to the terminal server to support both AppleTalk Remote Access Protocol (ARAP) and Defender authentication through terminal server. Previously, MAX products supported Defender authentication through the terminal server, but supported ARA through the ARA parameter at Ethernet > Answer > Encaps.

The newer parameter, also called ARA, has been added to the following location:

Ethernet > Mod Config > TServ

Before you use Defender for ARAP

Before you use Defender authentication through the terminal server, you must:

- Have ARA 2.0 or 2.1.

Note: Due to a Defender limitation, ARA 3.0 and above does not support this new feature.

- Use the reserved name Defender with a blank password in your ARA client software.
- Have Defender security software.
- Have the following scripts:
 - Direct-to-Defender authentication, ASCDIR.SCR
 - RADIUS-to-Defender authentication, ASCAC.SCR.

These scripts are available on the World Wide Web at <http://www.ascend.com> or you can contact Ascend Customer Service.

- If you will be using RADIUS to Defender authentication, you must add Authentication-Type = DEFENDER to the RADIUS users file.

ARA

Description: Specifies whether the MAX supports ARAP and Defender authentication through terminal server.

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

- Yes enables the MAX terminal server to support ARAP and Defender authentication, provided they meet all other connection criteria.
- No means the MAX will support neither ARAP nor Defender authentication through terminal server. You must change this setting regardless of whether you want to use direct-to-Defender or RADIUS-to-Defender authentication.

Example: ARA=Yes

Dependencies: If you have set AppleTalk or terminal server (TServ) No, ARA becomes not applicable (N/A).

Location: Ethernet > Mod Config > TServ

See Also: AppleTalk, Encaps, ARA Parameter, RADIUS Configuration Guide

Dial-in PPP support for AppleTalk

You can configure an Ascend unit so that individual users can dial into an AppleTalk network using a PPP dialer, such as AppleTalk Remote Access 3.0 and Pacer PPP. The MAX does not need to be set up as an AppleTalk router to support dial-in PPP to AppleTalk.

Overview

The following changes have been made to the user interface to support this feature:

- A new parameter, Peer, has been added to the Connection profile
- A new menu, AppleTalk Options, has been added to the Answer profile.

System requirements

This feature is supported only in MAX system loads that support AppleTalk routing. These loads are indicated by an “a” in the filename (for example, tba.m4, ba.m2).

Note: Many of the MAX updates with AppleTalk routing are “fat loads.” For more details refer to the release note entitled “Larger executable load images (“fat loads”) enabled.”

Configuring dial-in PPP for AppleTalk

You can set up a MAX to allow an AppleTalk client to dial in using PPP in two ways:

- using a Connection profile
- using a Name/Password profile

Configuring an AppleTalk PPP connection using a Connection profile

- 1 Open the Ethernet > Mod Config menu.
- 2 Set Appletalk=Yes.
- 3 Open the appropriate Connection profile.
- 4 Set Route Appletalk=Yes.
- 5 Open the AppleTalk options menu.

```
90-103 apple
  AppleTalk options...
    Peer=Dialin
    Zone Name=N/A
    Net Start=N/A
    Net End=N/A
```

- 6 Set the Peer parameter to indicate whether the connection for this profile is a single user PPP connection or a router

Peer=Dialin indicates that the profile is for a single user PPP connection. All other fields in the AppleTalk options menu are N/A. If you select Peer=Dialin, you have completed the configuration; close the AppleTalk Options menu and save your changes.

Peer=Router indicates that the profile is for a connection with a router (such as an Ascend Pipeline unit). If you select Peer=Router, you will need to configure the other fields in the AppleTalk options menu by continuing with through step 11

Note: Peer=Router works the same way that AppleTalk routing worked before this feature. The following steps are given here for convenience, and duplicate the existing documentation for AppleTalk routing.

- 7 Configure the AppleTalk zone name for the Ascend unit in the AppleTalk options submenu of the Ethernet Configuration Profile.

If there are other AppleTalk routers on the network, you must configure the zone names and network ranges to coincide with the other routers on the LAN.

The default for the Zone Name field is blank. Enter up to 33 alphanumeric characters to identify the zone name for the unit you are configuring.

Note: These fields will all display N/A if you have not enabled AppleTalk in the Ethernet Mod Config menu.

- 8 Specify whether the Ascend unit will be a seed or non-seed router. The default value for AppleTalk Router is Off.
 - You assign the network range and zone name configuration for a seed router. There must be at least one seed router on a routed AppleTalk network. Select AppleTalk Router=Seed for this option.
 - A non-seed router learns network number and zone information from other routers. Select AppleTalk Router=Non-Seed for this option.

If you choose Non Seed or Off, then Net Start, Net End, Default Zone, and Zone Name #x are N/A.

If you are configuring a non-seed router and are using Names/Passwords, go to Configuring an AppleTalk PPP connection using a Name/Password profile.

- 9 If you are configuring the Ascend unit as a seed router, specify the network range for the network to which the Ascend unit is attached.

Net Start and Net End define the network range for nodes attached to this network. Valid entries for these fields are in the range from 1 to 65199. If there are other AppleTalk routers on the network, you must configure the network ranges to coincide with the other routers.
- 10 Specify the default zone name for nodes on the Ascend unit's internet.

Enter up to 33 alphanumeric characters for the default zone name. The default for this field is blank.

The default zone is the one used by a node in the network for which you are configuring the Connection Profile until another zone name is explicitly selected by the node.
- 11 Specify the zone names that the platform can seed.

The Pipeline can seed up to 5 zones, and the MAX can seed up to 32. Enter up to 33 alphanumeric characters in zone name fields.

Configuring an AppleTalk PPP connection using a Name/Password profile

- 1 Open the Ethernet > Mod Config menu.
- 2 Set Appletalk=Yes.
- 3 Open the PPP Options menu of the Answer profile.
- 4 Set Route Appletalk=Yes.

- 5 Open the Appletalk options submenu of the PPP options menu.

```
90-103 apple
      AppleTalk options...
      Peer=Dialin
```

- 6 Set the Peer parameter to indicate whether the connection for this profile is a single user PPP connection or a router

Peer=Dialin indicates that the profile is for a single user PPP connection. All other fields in the AppleTalk options menu are N/A. If you select Peer=Dialin, you have completed the configuration; close the AppleTalk Options menu and save your changes.

Peer=Router indicates that the profile is for a connection with a router (such as an Ascend Pipeline unit). If you select Peer=Router, you will need to configure the other fields in the AppleTalk options menu by continuing with Step 7 through Step 11 in Configuring an AppleTalk PPP connection using a Name/Password profile.

Note: Peer=Router works the same way that AppleTalk routing worked before this feature. The Step 7 through Step 11 are given here for convenience, and duplicate the existing documentation for AppleTalk routing.

SecurID authentication for AppleTalk Remote Access (ARA) users

This release enables ARA callers to use AppSecurID authentication by contacting an ACE server through a Connection profile, Password profile, or RADIUS user profile. Previously, ARA callers could use only direct name and password authentication without the use of an external authentication server.

How an ARA caller uses SecurID authentication

An ARA caller can use SecurID authentication in any of the following ways:

- Using a Connection profile
- Using a Password profile
- Using a RADIUS user profile

If the user has a RADIUS user profile, he or she must have the username "SecurID".

For information on setting up a profile to contact an external authentication server, including an ACE server, see the *MAX Security Supplement*.

Once the user makes the initial connection, SecurID authentication begins with a pop-up screen on the Macintosh. At this point, the user must enter the "User ID" and "Passcode". If the user enters incorrect values, he or she gets two more tries to authenticate before the connection fails.

If the user is required to enter a new PIN, a pop-up screen prompts for this information. The user has three chances to enter the correct PIN. Once the new PIN is accepted, a pop-up screen instructs the Macintosh user to wait for the token code to change and then to log in with the new PIN and token code.

The SecurID client module must be version 1.3 or later.

Multiband features

Enable data-service calls for the Multiband MAX 1800, 2000, and 4000

Data-service functionality is now turned on for Multiband MAX units.

The Multiband MAX products have AIM cards for connecting a video system to a network for videoconferencing. The video systems do not use analog links. They require digital connections. Therefore, Ascend has enabled data-service calls on the Multiband MAX products. The entry `DataCall Installed` appears in the System Options status window, indicating that data service calls are enabled.

Suppress the display of the second T1 line on the Multiband MAX 2000

The Multiband MAX 2000 is a new version of the MAX 2000. This version supports only a single T1 PRI line. Therefore, the Net/T1 window does not contain an entry for a second line.

MAX 2000 main Status menu

The MAX 2000 main Status menu looks like this one:

```
10-000 Net/T1
  10-100 Line 1 Stat
  10-200 Line error
  10-300 net Options
```

Suppress the T1-CSU and Serial WAN menus for the Multiband MAX 1800, 4002, and 4004

The Multiband MAX is a new version of the MAX. This version supports only switched T1 PRI lines, not switched and nailed-up T1 PRI lines as in the MAX product. In addition, the Multiband MAX product does not include a Serial WAN option. Therefore, the T1-CSU and Serial WAN menu items do not appear on the Multiband MAX.

Multiband MAX 1800 main Edit menu

The Multiband MAX 1800 main Edit menu looks like this one:

```
Main Edit Menu
  00-000 System
  10-000 Net/BRI
  20-000 Empty
  30-000 Empty
```

Multiband features

Suppress the T1-CSU and Serial WAN menus for the Multiband MAX 2000

```
40-000 Reserved
50-000 Ethernet
```

Multiband MAX 4002 and 4004 main Edit menu

The Multiband MAX 4002 and 4004 main Edit menu looks like this one:

```
Main Edit Menu
00-000 System
10-000 Net/T1
20-000 Net/T1
30-000 Empty
40-000 Empty
50-000 Empty
60-000 Empty
70-000 Empty
80-000 Empty
90-000 Ethernet
A0-000 Ether Data
B0-000 Reserved
```

Note: The second Net/T1 line appears on the Multiband MAX 4004 only.

Suppress the T1-CSU and Serial WAN menus for the Multiband MAX 2000

The Multiband MAX 2000 is a new version of the MAX 2000. This version supports only a switched T1 PRI line, not switched and nailed-up lines as in the MAX 2000 product. In addition, the Multiband MAX product does not include a Serial WAN option. Therefore, the T1-CSU and Serial WAN menu items do not appear on the Multiband MAX.

Multiband MAX 2000 main Edit menu

The Multiband MAX 2000 main Edit menu looks like this one:

```
Main Edit Menu
00-System
10-Net/T1
20-Empty
30-Empty
40-Reserved
50-Ethernet
```

Multiband simulation restricts MAX functionality

When you enable Multiband simulation, bridging and IP routing are disabled, and you cannot use certain terminal server commands.

Features disabled

Using the Multiband simulation functionality has the following effects:

- Bridging is disabled.
You cannot set the Bridging parameter to Yes. If you attempt to do so, this error message appears:

```
#226 Multiband does not permit Bridging. Set Bridging=No.
```
- IP routing is disabled.
IP packets are not routed through the MAX.
- You cannot use these terminal server commands:
 - close
 - ipxping
 - open
 - resume
 - rlogin
 - telnet

Multiband only option added

An option has been added to the platform hash codes that restricts the MAX to multiband functions only. With this hash code installed, the word "Multiband" is displayed in the POST screen and Sys Options status display. Previously, the word "multiband" was always displayed.

Multiband-only option

When the multiband-only hash code is installed, the MAX displays the word "Multiband" in the POST screen and in the Sys Options window. The MAX functions are restricted to multiband-only features.

\Multiband functionality includes AIM, BONDING, and other functions associated with the MAX's Host/Dual and Host/6 slot cards. Essentially, multiband functionality includes options in the Host/Dual, Host/Quad, or Host/6 (AIM-6) branches of the edit menu. Specifically, Host Interface Profiles, Port Profiles, Call Profiles, and Port Diag commands.

Sys Options menu addition

Description: Installed options appear in the Sys Options status menu. The multiband only options. If the Multiband-only hash option is installed, the Sys Options status window will contain "Multiband MAX."