



MAX TNT™ True Access™ Operating System (TAOS)

8.0.0 Release Note

Copyright© 2000 Lucent Technologies. All Rights Reserved.

This material is protected by the copyright laws of the United States and other countries. It may not be reproduced, distributed, or altered in any fashion by any entity (either internal or external to Lucent Technologies), except in accordance with applicable agreements, contracts or licensing, without the express written consent of Lucent Technologies, Inc.

Notice

Every effort was made to ensure that the information in this document was complete and accurate at the time of printing. However, information is subject to change.

Security Statement

In rare instances, unauthorized individuals make connections to the telecommunications network through the use of access features.

Trademarks

DSLPipe, DSLMAX, DSL Terminator, DSLTNT, MAX, MAX TNT, MultiDSL, MultiVoice, Pipeline, GRF, NavisRadius, NavisAccess, and Stinger are trademarks of Lucent Technologies. Other trademarks and trade names mentioned in this publication belong to their respective owners.

Ordering Information

To order copies of this document, contact your Lucent Technologies representative or reseller.

Support Telephone Numbers

For a menu of support and other services, call (800) 272-3634. Or call (510) 769-6001 for an operator.

Customer Service

Customer Service provides a variety of options for obtaining information about Lucent products and services, software upgrades, and technical assistance.

Finding information and software on the Internet

Visit the Web site at <http://www.ascend.com> for technical information, product information, and descriptions of available services.

Visit the FTP site at <ftp.ascend.com> for software upgrades, release notes, and addenda to this manual.

Obtaining technical assistance

You can obtain technical assistance by telephone, email, fax, modem, or regular mail, as well as over the Internet.

Enabling Lucent to assist you

If you need to contact Lucent for help with a problem, make sure that you have the following information when you call or that you include it in your correspondence:

- Product name and model.
- Software and hardware options.
- Software version.
- If supplied by your carrier, Service Profile Identifiers (SPIDs) associated with your line.
- 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 Lucent product.
- Type of computer you are using.
- Description of the problem.

Calling Lucent from within the United States

In the U.S., you can take advantage of Priority Technical Assistance or an Advantage service contract, or you can call to request assistance.

Priority Technical Assistance

If you need to talk to an engineer right away, call (900) 555-2763 to reach the Priority Call queue. The charge of \$2.95 per minute does not begin to accrue until you are connected to an engineer. Average wait times are less than three minutes.

Advantage Services

Advantage Services is a comprehensive selection of services. Installation services help get your Lucent Wide Area Network (WAN) off to the right start. Ongoing maintenance and

support services provide hardware and software solutions to keep your network operating at peak performance. For more information, call (800) 272-3634, or access the Web site at www.ascend.com and select Services and Support, then Advantage Services.

Other telephone numbers

For a menu of Lucent's services, call (800) 272-363. Or call (510) 769-6001 for an operator.

Calling Lucent from outside the United States

You can contact Lucent by telephone from outside the United States at one of the following numbers:

Telephone outside the United States	(510) 769-8027
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

For the Asia Pacific Region, you can find additional support resources at <http://apac.ascend.com>

Obtaining assistance through correspondence

Lucent maintains two email addresses for technical support questions. One is for customers in the United States, and the other is for customers in Europe, the Middle East, and Asia. If you prefer to correspond by fax, BBS, or regular mail, please direct your inquiry to Lucent's U.S. offices. Following are the ways in which you can reach Customer Service:

- Email from within the U.S.—support@ascend.com
- Email from Europe, the Middle East, or Asia—EMEAsupport@ascend.com
- Fax—(510) 814-2312
- Customer Support BBS (by modem)—(510) 814-2302

Write to Lucent at the following address:

Attn: Customer Service
Lucent Technologies
1701 Harbor Bay Parkway
Alameda, CA 94502-3002

Table of Contents

Customer Service	3
What's new?	13
How to use this release note.....	13
TAOS built-in features	13
TAOS extensions	16
Notices and caveats	18
Notice of modified RADIUS port and ID space defaults	18
Notice of modified behavior during IPCP negotiation	19
Notice of discontinuance of software support	19
Notice of discontinuance of MAX TNT support for DSL.....	19
Notice of deprecated management features	20
Notice about upgrading slot cards	20
Caveats in this release	20
Upgrade and downgrade procedures	23
Requirements and recommendations	23
Obtaining the MAX TNT TAOS 8.0.0 software	23
Local access to the unit recommended	23
32-MB JEDEC DRAM card required for this release	24
Flash size limitations for this upgrade	24
Saving the system configuration.....	25
Upgrade instructions	25
Upgrading a standalone MAX TNT unit	25
Upgrading a multishelf MAX TNT unit.....	26
Downgrade instructions	28
Downgrading a standalone MAX TNT unit	28
Downgrading a multishelf MAX TNT unit	29
Built-in features in MAX TNT TAOS 8.0.0.....	31
Modem Manager.....	31
Firmware versions for digital modems	31
Firmware versions for MultiDSP cards	31
Series56 III modem card support.....	31
Expanded MultiDSP card support	32
Stac compression for asynchronous calls	32
WAN Access Server	34
Ethernet-3 slot card support.....	34
OC3-ATM slot card support.....	34
HDLC2-EC-C slot card support	35

STM-0 slot card support	35
Bandwidth Allocation Control Protocol (BACP)	35
X.75 frame size increase to 2048 bytes	37
Support for multidigit trunk groups	37
MAX TNT stacking	39
Authentication, Authorization, and Accounting (AAA)	54
Sharing profiles on a per-user basis	54
New settings for CLID-Auth-Mode	55
Extensions to DNIS management	55
Bidirectional CHAP	58
Username available for first-tier DNIS authentication	68
NAS-Port-Type specification for local profiles	69
RADIUS: Enhanced IETF compliance in VSA compatibility mode	70
RADIUS: Tunnel attribute sets with tags and preferences	70
RADIUS: Support for MS-CHAP authentication	73
RADIUS: Authentication of Telnet sessions	73
RADIUS: Inclusion of data and transmit rates in access request	74
RADIUS: TCP-Clear-Quiet value added to Login-Service attribute	75
RADIUS: Toggle boot requests	76
RADIUS: Overriding the Answer-Defaults authentication method	76
ISDN callback	77
Management Agent	87
DOS-compatible FAT-16 flash memory format	87
Additional onboard memory for extended profiling	90
Pattern detection and reporting in the TCP-Clear data stream	93
Telnet access control list	95
Periodic log message for reporting the software version	96
Remote management of other units	96
Command for displaying filter information	98
Customizable screen width up to 255 characters	101
TCP-Clear login host IP address reported	102
Coredump generated by specific Warning messages	103
Event logging when an operator downs or resets a slot card	104
Support for multiple Syslog servers	106
Full backup queue type now identified in Syslog	107
Grep-like capability added to certain commands	108
CLID display option for Userstat command	110
Options for displaying the system log	111
Option for displaying line status	112
Dynamic remote filters	113
MaxTap support	119
Userstat options to display address and username	119
Selectable call-logging server	121
Enhanced reporting of HDLC errors	122
SNMP: Partial support for ATM MIB RFC 2515	123
SNMP: Trap for dropped call-logging packets	124
SNMP: Configuration of TFTP port for network management	124
SNMP: New Config-Change trap	125
SNMP: Trap for console state change now displays IP address	126
SNMP: Support for loopback modes in DS1 MIB	126
SNMP: Fatal log table	126
SNMP: Idle Time variable in the active session table	127

SNMP: Variables added from event.mib to call.mib.....	127
SNMP: New MIB variables for summarizing B channel states	128
SNMP: Modified method for adding SNMP object IDs.....	128
SNMP: Details on terminating access resources	128
SNMP: Limited support for RFC 2574 user-based security model.....	130
IP routing.....	133
OSPF nonbroadcast multiaccess (NBMA) support	133
OSPF area border router (ABR) support	137
OSPF MD5 16-byte authentication key	137
RFC 1850 OSPF traps.....	138
OSPF reconfiguration restart no longer required.....	141
Per-connection Microsoft WINS assignment	142
Private routing tables	146
Port redirection	153
IP pool chaining	155
Don't Fragment option added to Ping command.....	161
Commands for monitoring multipath route caching.....	161
Store-and-forward IP fax	163
Incoming and outgoing IP faxes	163
System parameters for IP fax modem usage.....	164
Configuring the MAX TNT for IP fax.....	166
Example of an IP fax configuration for incoming faxes.....	168
Example of an IP fax configuration for outgoing faxes.....	169
Fax hangup codes and disconnect cause codes.....	171
IP fax call accounting	171
SNMP changes for IP fax operation	171
RADIUS changes for IP fax operation	172
Syslog changes for IP fax operation	174
Redialer support on MultiDSP card for store-and-forward fax	174
Extension features in MAX TNT TAOS 8.0.0	175
Global Digital Access	175
PHS Internet Access Forum Standard (PIAFS) 2.1 on MultiDSP.....	175
Asynchronous V.110 on MultiDSP	175
R2 CLID processing for New Zealand	176
R2 CLID processing for Thailand	176
R2 CLID processing for Israel.....	177
R2 CLID processing for Mexico	177
R2 signaling for Kuwait.....	178
Support for ISDN network-side emulation (T1 and E1).....	178
ISDN NFAS support for Japanese switch types	179
Frame Relay PVCs over switched ISDN connections.....	179
Frame Relay switched virtual circuits (SVCs)	182
Multilink Frame Relay (MFR).....	187
DLCI bundling in MFR	198
MFR circuit switching	204
Configurable VPI-VCI ranges	207
OAM loopback for DS3 ATM PVC fault management	208
ATM-direct	209
ATM switched virtual circuits (SVCs)	213
ATM traffic shaping	225

ATM-Frame Relay transparent-mode circuits (FRF.8)	227
ATM-Frame Relay virtual channel trunking	230
Signaling System 7 (SS7)	233
System requirements for SS7 operations	233
MAX TNT as terminator of data calls in an SS7 network.....	234
MAX TNT as terminator of voice and data calls in an SS7 network	234
Interface between a signaling gateway and MAX TNT	236
Incoming calls	236
Continuity tests	236
SS7-Gateway profile settings.....	236
T1 lines as SS7 data trunks	240
E1 lines as SS7 data trunks	241
V.110 bearer capability for SS7 calls using IPDC.....	242
SS7 link establishment timer	242
Support for 2-wire continuity check on T1 lines	243
Outgoing continuity tests on T1 and T3	244
Digital milliwatt tone support on T1 and T3	244
Analog milliwatt tone and variable tone support.....	245
IPDC enhancements for reporting VoIP call statistics	245
Statistics and error reporting on SS7 connections	247
Cause codes for SS7 ASGCP calls to the MAX TNT	250
SS7 IPDC support for call ID and disconnect cause codes	251
SNMP: Support for the SS7 MIB (ascend 29).....	254
SNMP: Support for SS7 link-state trap.....	254
SNMP: Idle time attribute in active session table.....	254
MultiVoice operations.....	255
System requirements for VoIP.....	255
Ethernet requirements for VoIP processing.....	256
VoIP call routing.....	256
Overview of VoIP in an H.323v2 environment	257
Overview of VoIP in an SS7 IPDC 0.12 environment	258
General system configuration for VoIP support	259
VoIP call management and performance settings.....	262
Real-time fax (T.38)	262
Transparent modem	264
Using transparent modem with real-time fax	264
New VoIP profile settings in MAX TNT TAOS 8.0.0.....	265
IPDC message support for modifying parameters	269
IPDC message support for T.38 fax and transparent modem	270
New trunk features for VoIP calls	270
RT-24 (proprietary) codec support	273
G.728 codec support	273
SNMP: Support for the VoIP MIB (ascend 28).....	274
SNMP: Traps for VoIP-related conditions	275
NavisAccess support for VoIP call reporting	276
Tunneling	278
IPSec for L2TP tunnels and TCP-Clear.....	278
L2TP timer options	288
L2TP list attempts.....	289
Multiple endpoints for tunnel sessions (RADIUS only).....	290
Secondary tunnel server for L2TP and L2F tunnels (local profiles)	290
Optional L2TP system name.....	293

Limited support for Layer 2 Forwarding (L2F).....	293
Virtual routing.....	301
Virtual router (VRouter) DNS	301
VRouter support for L2TP connections.....	303
Short-duration transaction network (SDTN).....	304
SDTN operation.....	304
Transaction-Server profiles.....	305
Dial-in connections for transaction clients	308
Corrections in MAX TNT TAOS 8.0.0.....	315

Figures

Figure 1.	Stacked bundle consisting of three links.....	39
Figure 2.	Stacking control messages to establish a bundle	40
Figure 3.	Stacking data and control packets.....	41
Figure 4.	Stacking data interfaces in star configuration.....	42
Figure 5.	Stacking MAX TNT and MAX units together	42
Figure 6.	Stacked peers authenticating calls via RADIUS.....	45
Figure 7.	Example stack of two MAX TNT units.....	48
Figure 8.	Example stack of a MAX TNT and MAX unit.....	51
Figure 9.	Bidirectional CHAP for all incoming calls to the MAX TNT unit	58
Figure 10.	Bidirectional CHAP for selected calls	60
Figure 11.	Bidirectional CHAP in a multiprovider network	66
Figure 12.	OSPF nonbroadcast multiaccess (NBMA) network	135
Figure 13.	Port redirection to an HTTP server.....	155
Figure 14.	Incoming IP fax from fax machine to Internet.....	163
Figure 15.	Outgoing IP fax from Internet to fax machine.....	164
Figure 16.	Receiving and forwarding incoming IP faxes.....	168
Figure 17.	Sending an outgoing IP fax to a fax machine	170
Figure 18.	Switched PVC to a Frame Relay switch	180
Figure 19.	Terminating SVC on a Frame Relay interface.....	182
Figure 20.	Dial-out SVC on a Frame Relay interface	183
Figure 21.	SVC between MAX TNT units with an intervening Frame Relay switch	186
Figure 22.	MFR DTE-DTE aggregation	188
Figure 23.	MFR peers with three data links supporting two DLCIs	189
Figure 24.	Sample MFR configuration	192
Figure 25.	Example of MFR on per-DLCI basis.....	199
Figure 26.	MFR circuit switching from a bundle to a single T1 interface	204
Figure 27.	16-bit VPI-VCI range	207
Figure 28.	ATM permanent virtual circuit	208
Figure 29.	ATM-direct concentrating PPP calls to an ATM interface.....	211
Figure 30.	Terminating SVC on an ATM interface	213
Figure 31.	Dial-out SVC on an ATM interface.....	214
Figure 32.	Subfields in the AESA address formats.....	215
Figure 33.	Example ATM SVC with DCC-AESA addresses	222
Figure 34.	Example traffic shaping setup.....	226
Figure 35.	ATM-Frame Relay circuit.....	228
Figure 36.	N:1 circuit between multiple Frame Relay hosts and an ATM trunk	230
Figure 37.	Circuit using virtual channel trunking	232
Figure 38.	MAX TNT terminating data calls in an SS7 network.....	234
Figure 39.	MAX TNT terminating voice and data calls in an SS7 network	235
Figure 40.	Simplified view of VoIP call routing within the MAX TNT	257
Figure 41.	Example diagram of MultiVoice in H.323 environment	258
Figure 42.	Example diagram of MultiVoice in SS7 environment.....	259
Figure 43.	Secure IPsec L2TP tunneling configuration	280

Figure 44.	IPsec tunnel mode for TCP-Clear between gateways.....	281
Figure 45.	IPsec transport mode for TCP-Clear with dial-in host.....	282
Figure 46.	Primary and secondary L2TP tunnel endpoints	291
Figure 47.	L2F tunneling.....	294
Figure 48.	L2TP tunnels built on separate VRouters	303
Figure 49.	Sample SDTN setup.....	305
Figure 50.	Transaction servers with redundant Ethernet connections.....	308

What's new?

The True Access™ Operating System (TAOS) contains a foundation of built-in software features for WAN access environments, as well as optional extensions that require separate licensing to support a wide variety of WAN access environments.

This release note describes all new features and extensions that have been introduced for MAX TNT™ units since TAOS 7.0.0. Some of the features have been introduced in earlier 7.x releases, as indicated in the tables below.

How to use this release note

The Table of Contents and the tables in the following section list the TAOS features in this release. If you are reading this release note in PDF format, you can click the feature name to go directly to the Feature.

For information about obtaining the software described in this release note, see “Obtaining the MAX TNT TAOS 8.0.0 software” on page 23.

TAOS built-in features

Table 1 through Table 6 show built-in features that have been added to TAOS since the last major release. Features that were also available in earlier 7.x releases are indicated.

Table 1. TAOS 8.0.0 Modem Manager features

Feature	New	In 7.x
Firmware versions for digital modems	√	
Firmware versions for MultiDSP cards	√	
Series56 III modem card support		√
Expanded MultiDSP card support	√	
Stac compression for asynchronous calls		√

Table 2. TAOS 8.0.0 WAN Access Server features

Feature	New	In 7.x
Ethernet-3 slot card support		√
OC3-ATM slot card support		√
HDLC2-EC-C slot card support		√
STM-0 slot card support	√	
Bandwidth Allocation Control Protocol (BACP)		√
X.75 frame size increase to 2048 bytes		√
Support for multidigit trunk groups		√
MAX TNT stacking		√

Table 3. TAOS 8.0.0 AAA Server features

Feature	New	In 7.x
Sharing profiles on a per-user basis	√	
New settings for CLID-Auth-Mode		√
Extensions to DNIS management		√
Bidirectional CHAP	√	
Username available for first-tier DNIS authentication	√	
NAS-Port-Type specification for local profiles		√
RADIUS: Enhanced IETF compliance in VSA compatibility mode	√	
RADIUS: Tunnel attribute sets with tags and preferences	√	
RADIUS: Support for MS-CHAP authentication	√	
RADIUS: Authentication of Telnet sessions		√
RADIUS: Inclusion of data and transmit rates in access request		√
RADIUS: TCP-Clear-Quiet value added to Login-Service attribute		√
RADIUS: Toggle boot requests	√	
RADIUS: Overriding the Answer-Defaults authentication method	√	
ISDN callback		√

Table 4. TAOS 8.0.0 Management Agent features

Feature	New	In 7.x
DOS-compatible FAT-16 flash memory format	√	
Additional onboard memory for extended profiling		√
Pattern detection and reporting in the TCP-Clear data stream		√
Telnet access control list	√	
Periodic log message for reporting the software version		√
Remote management of other units		√
Command for displaying filter information		√
Customizable screen width up to 255 characters		√
TCP-Clear login host IP address reported		√
Coredump generated by specific Warning messages	√	
Event logging when an operator downs or resets a slot card	√	
Support for multiple Syslog servers	√	
Full backup queue type now identified in Syslog	√	
Grep-like capability added to certain commands	√	
CLID display option for Userstat command	√	
Options for displaying the system log	√	
Option for displaying line status	√	
Dynamic remote filters	√	
MaxTap support	√	

Table 4. TAOS 8.0.0 Management Agent features (continued)

Feature	New	In 7.x
Userstat options to display address and username	√	
Selectable call-logging server	√	
SNMP: Partial support for ATM MIB RFC 2515	√	
SNMP: Trap for dropped call-logging packets		√
SNMP: Configuration of TFTP port for network management		√
SNMP: New Config-Change trap		√
SNMP: Trap for console state change now displays IP address	√	
SNMP: Support for loopback modes in DS1 MIB	√	
SNMP: Fatal log table	√	
SNMP: Idle Time variable in the active session table	√	
SNMP: Variables added from event.mib to call.mib	√	
SNMP: New MIB variables for summarizing B channel states		√
SNMP: Modified method for adding SNMP object IDs		√
SNMP: Details on terminating access resources		√
SNMP: Limited support for RFC 2574 user-based security model	√	

Table 5. TAOS 8.0.0 IP Router features

Feature	New	In 7.x
OSPF nonbroadcast multiaccess (NBMA) support	√	
OSPF area border router (ABR) support	√	
OSPF MD5 16-byte authentication key	√	
RFC 1850 OSPF traps	√	
OSPF reconfiguration restart no longer required	√	
Per-connection Microsoft WINS assignment	√	
Private routing tables	√	
Port redirection	√	
Port redirection		√
IP pool chaining		√
Don't Fragment option added to Ping command		√
Commands for monitoring multipath route caching		√

Table 6. TAOS 8.0.0 IP fax features

Feature	New	In 7.x
Store-and-forward IP fax		√
IP fax call accounting		√
Redialer support on MultiDSP card for store-and-forward fax	√	

TAOS extensions

Table 7 through Table 12 show extension features that have been added to TAOS since the last major release. Extension features are available when the appropriate software license has been enabled. Features that were also available in earlier 7.x releases are indicated.

Table 7. TAOS 8.0.0 Global Digital Access extension features

Feature	New	In 7.x
PHS Internet Access Forum Standard (PIAFS) 2.1 on MultiDSP	√	
Asynchronous V.110 on MultiDSP		√
R2 CLID processing for New Zealand	√	
R2 CLID processing for Thailand	√	
R2 CLID processing for Israel	√	
R2 CLID processing for Mexico	√	
R2 signaling for Kuwait		√
Support for ISDN network-side emulation (T1 and E1)	√	
ISDN NFAS support for Japanese switch types		√
Frame Relay PVCs over switched ISDN connections	√	
Frame Relay switched virtual circuits (SVCs)	√	
Multilink Frame Relay (MFR)		√
DLCI bundling in MFR		√
MFR circuit switching	√	
Configurable VPI-VCI ranges		√
OAM loopback for DS3 ATM PVC fault management	√	
ATM-direct	√	
ATM switched virtual circuits (SVCs)	√	
ATM traffic shaping	√	
ATM-Frame Relay transparent-mode circuits (FRF.8)	√	
ATM-Frame Relay virtual channel trunking	√	

Table 8. TAOS 8.0.0 Signaling System 7 (SS7) extension features

Feature	New	In 7.x
System requirements for SS7 operations	√	
MAX TNT as terminator of data calls in an SS7 network		√
MAX TNT as terminator of voice and data calls in an SS7 network	√	
SS7-Gateway profile settings	√	
T1 lines as SS7 data trunks		√
E1 lines as SS7 data trunks		√
V.110 bearer capability for SS7 calls using IPDC	√	
SS7 link establishment timer	√	
Support for 2-wire continuity check on T1 lines	√	

Table 8. TAOS 8.0.0 Signaling System 7 (SS7) extension features (continued)

Feature	New	In 7.x
Outgoing continuity tests on T1 and T3	√	
Digital milliwatt tone support on T1 and T3	√	
Analog milliwatt tone and variable tone support	√	
IPDC enhancements for reporting VoIP call statistics	√	
Statistics and error reporting on SS7 connections	√	
Cause codes for SS7 ASGCP calls to the MAX TNT	√	
SS7 IPDC support for call ID and disconnect cause codes	√	
SNMP: Support for the SS7 MIB (ascend 29)	√	
SNMP: Support for SS7 link-state trap	√	
SNMP: Idle time attribute in active session table	√	

Table 9. TAOS 8.0.0 MultiVoice™ extension features

Feature	New	In 7.x
System requirements for VoIP	√	
Ethernet requirements for VoIP processing		√
VoIP call routing		√
General system configuration for VoIP support		√
VoIP call management and performance settings		√
Real-time fax (T.38)		√
Transparent modem		√
Storing voice announcements in the FAT-16 flash memory file system	√	
Allowing fallback to alternate codecs	√	
Deactivating trunks used for VoIP calls	√	
Enabling early ringback	√	
Trunk prefixing	√	
IPDC message support for modifying parameters	√	
Configurable interdigit timer for T1 inband signaling		√
Delaying charges until call is answered (true connect)		√
Gatekeeper CLID substitution		√
RT-24 (proprietary) codec support	√	
G.728 codec support	√	
SNMP: Support for the VoIP MIB (ascend 28)	√	
SNMP: Traps for VoIP-related conditions	√	
NavisAccess support for VoIP call reporting	√	

Table 10. TAOS 8.0.0 Tunneling extension features

Feature	New	In 7.x
IPSec for L2TP tunnels and TCP-Clear	√	
L2TP timer options	√	

Table 10. TAOS 8.0.0 Tunneling extension features (continued)

Feature	New	In 7.x
L2TP list attempts	√	
Multiple endpoints for tunnel sessions (RADIUS only)	√	
Secondary tunnel server for L2TP and L2F tunnels (local profiles)	√	
Optional L2TP system name		√
Limited support for Layer 2 Forwarding (L2F)	√	

Table 11. TAOS 8.0.0 Virtual Routing extension features

Feature	New	In 7.x
Virtual router (VRouter) DNS	√	
VRouter support for L2TP connections	√	

Table 12. TAOS 8.0.0 Short duration transaction network extension

Feature	New	In 7.x
Short-duration transaction network (SDTN)	√	

Notices and caveats

Notice of modified RADIUS port and ID space defaults

Note: This modification could cause authentication failures with RADIUS servers that do not support distinct UDP source ports. If your RADIUS server does not support authentication requests from multiple source ports, you must reset the modified parameters to their previous values.

The default settings for User Datagram Protocol (UDP) source ports and ID spaces for communication with a RADIUS server have been changed from single to multiple. Following are the relevant parameters, shown with the new default settings:

```
[EXTERNAL-AUTH]
rad-id-space = distinct
rad-id-source-unique = port-unique
```

MAX TNT units can use either a single global source UDP port for all slot cards, or a unique port for each card. Similarly, a unit can use one ID space for both authentication and accounting requests, or a distinct space for each type of request.

Previous TAOS versions recommended the use of multiple source ports and ID spaces for performance reasons, and because use of a single source port and ID space reduces the number of simultaneous requests that the unit can generate. However, the default settings configured a single global source port and ID space to ensure compatibility with all RADIUS servers.

In this release, the default settings have been changed to the recommended values.

If the system was already using the recommended settings, this change will have no effect.

Systems that used the previous default settings will respond as follows:

- If the RADIUS server supports distinct source ports, the system will experience a slight improvement in performance.
- If the RADIUS server does not support distinct source ports, the system will experience problems with RADIUS authentication and accounting.

To communicate with RADIUS servers that do not support distinct source ports, you must modify the External-Auth profile as follows to restore the parameters to their previous values:

```
admin> read external-auth
EXTERNAL-AUTH read

admin> set rad-id-space = unified

admin> set rad-id-source-unique = system-unique

admin> write
EXTERNAL-AUTH written
```

Notice of modified behavior during IPCP negotiation

In previous releases, the MAX TNT unit's system address was used during IP Control Protocol (IPCP) negotiation. In previous releases, if the System-IP-Addr parameter was null, the shelf controller IP address was used.

With MAX TNT TAOS 8.0.0, the MAX TNT unit requires a valid System-IP-Addr setting to complete IPCP negotiation. For example, the following commands explicitly set the system address to the shelf controller IP address:

```
admin> get ip-int { {1 c 1} 0} ip-address
ip-address = 10.2.3.4

admin> read ip-global
IP-GLOBAL read

admin> set system-ip-addr = 10.2.3.4

admin> write
IP-GLOBAL written
```

Note: If the System-IP-Addr setting is null, the system terminates PPP connections during the IPCP negotiation phase.

Notice of discontinuance of software support

Software support has been discontinued for the MAX TNT Ethernet-0 slot card (TNT-SL-E10), the Fast (100 MB) Ethernet-1 slot card (TNT-SL-E100), and the older MAX TNT Hybrid Access slot cards (TNT-SL-HA128 and TNT-SL-HA192).

Notice of discontinuance of MAX TNT support for DSL

Support for digital subscriber loop (DSL) functionality is discontinued in MAX TNT units as of MAX TNT TAOS 8.0.0. The DSLTNT™ platform continues support for existing DSL products and will introduce additional DSL functionality in future releases.

Notice of deprecated management features

Use of the `if-admin` diagnostic command is deprecated. The functionality that was provided by the `-d` (down) and `-u` (up) options of the command is now provided by `read`, `set`, and `write` operations on one of the following profiles:

- The Admin-State-Perm-If profile for permanent interfaces such as a nailed interface
- The Admin-State-Phys-If profile for physical interfaces such as a T1 line

The other options of the `if-admin` command are not supported.

Use of the `call-log-radius-compat` parameter in the Call-Logging profile is deprecated in this software version.

The `callActiveIfIndex` and `callStatusIfIndex` objects in the call MIB are not supported in this software version.

The following objects are supported in this software version, but will not be supported in future software versions:

- The `lmodem.mib`
- The `resetStat` group in `ascend.mib`
- The `consoleTable`, `doTable`, and `hostStatusTable` in `ascend.mib`

Notice about upgrading slot cards

If you replace a MAX TNT Fast (100 MB) Ethernet-1 slot card (TNT-SL-E100) with a newer Ethernet card (TNT-SL-E10-100 or TNT-SL-E100-V-C), you must write new Ethernet profiles for the new card. The old Ethernet profiles do not carry forward.

If you replace an older MAX TNT Hybrid Access slot card (TNT-SL-HA128 or TNT-SL-HA192) with a newer Hybrid Access card (TNT-SL-HDLC2 or TNT-SL-HDLC2-EC-C), and if you replace a MAX TNT Series56 modem card (TNT-SL-48MOD-S56) with a newer Series56 card (TNT-SL-48MOD-S-C or TNT-SL-48MODV3-S-C), you must write new profiles for the new cards.

If you replace a Series56 modem card (TNT-SL-48MOD-S56, TNT-SL-48MOD-SGL, TNT-SL-48MOD-S-C or TNT-SL-48MODV3-S-C) with a MultiDSP card (TNT-SL-ADI-C, TNTV-SL-ADI-C, or APX8-SL-96DSP), you must write new profiles for the new cards.

For any slot whose card type is being changed, you should perform a `slot -r` command after downing (`slot -d`) or removing the existing card prior to inserting a new card type.

Caveats in this release

- Before changing an ATM connection's VPI-VCI assignment, you must disable the connection on a MAX TNT OC3 (Copper) ATM slot card (TNT-SL-OC3-C) or MAX TNT OC3 (Fibre) ATM slot card (TNT-SL-OC3-F).
- Multilink Protocol (MP) bonding of analog calls is supported, but some client modems and software may have compatibility problems.
- Configurable receive and transmit data rate limits are not supported on the MAX TNT unchannelized DS3-ATM slot card (TNT-SL-UDS3A). Configurable receive and transmit data rate limits *are* supported on the unchannelized DS3 Frame slot card (TNT-SL-UDS3).

- LAN-Modem profiles contain entries for 96 devices. For the 96-port MultiDSP card, all 96 entries in the profile are used. For 48-port modem cards (Series56 modem card (TNT-SL-48MOD-S56), Series56 II (TNT-SL-48MOD-S-C), and Series56 III (TNT-SL-48MODV3-S-C) cards), only the first 48 entries are used. For the 48-port MultiDSP card (TNTP-SL-ADI-C or TNTV-SL-ADI-C), every other entry in a LAN-Modem profile is used (odd ports only, from 1 to 95).

TAOS feature request IDs

The feature request IDs in Table 13 are included in this release.

Table 13. Feature Request IDs in this release

Feature ID	Description
51	Frame Relay switched virtual circuits (SVCs)
175	Frame Relay PVCs over switched ISDN connections
259742	SNMP: New Config-Change trap
259778	SNMP: Variables added from event.mib to call.mib
259937	Support for multiple Syslog servers
259938	TCP-Clear login host IP address reported
260023	Private routing tables
260147	SNMP: Details on terminating access resources
260257	Dynamic remote filters
260267	Bandwidth Allocation Control Protocol (BACP)
260342	L2TP list attempts
260346	Full backup queue type now identified in Syslog
260425	IPDC enhancements for reporting VoIP call statistics
260449	Grep-like capability added to certain commands
260598	SS7 IPDC support for call ID and disconnect cause codes
260600	Multiple endpoints for tunnel sessions (RADIUS only)
260614	IPDC message support for T.38 fax and transparent modem
501248	IPDC message support for T.38 fax and transparent modem
509674	SNMP: Limited support for RFC 2574 user-based security model
509682	RADIUS: Authentication of Telnet sessions
509702	ISDN callback
509718	Limited support for Layer 2 Forwarding (L2F)
509774	SNMP: New MIB variables for summarizing B channel states
509804	Per-connection Microsoft WINS assignment
509848	Telnet access control list
509854	Command for displaying filter information
509869	Secondary tunnel server for L2TP and L2F tunnels (local profiles)
509892	ATM-direct

Table 13. Feature Request IDs in this release (continued)

Feature ID	Description
509895	R2 CLID processing for Thailand
509898	ISDN NFAS support for Japanese switch types
509908	R2 signaling for Kuwait
509945	RADIUS: Toggle boot requests
510000	Userstat options to display address and username
510007	New settings for CLID-Auth-Mode
510027	RADIUS: Support for MS-CHAP authentication
510029	Virtual router (VRouter) DNS
510046	X.75 frame size increase to 2048 bytes
510054	Support for ISDN network-side emulation (T1 and E1)
510069	RADIUS: Tunnel attribute sets with tags and preferences
510148	Bidirectional CHAP

Upgrade and downgrade procedures

This section shows how to upgrade and downgrade the TAOS software of a MAX TNT unit.

Note: Digital subscriber loop (DSL) functionality is not supported in this release. See “Notice of discontinuance of MAX TNT support for DSL” on page 19.

Requirements and recommendations

These recommendations for upgrading MAX TNT units help ensure a smooth upgrade. If you must downgrade from this release to a previous one, please see “Downgrade instructions” on page 28.

Obtaining the MAX TNT TAOS 8.0.0 software

The MAX TNT TAOS 8.0.0 software consists of the following files:

Filename	Descriptions
<code>tntsrb.bin</code>	The boot loader. Both T1 and E1 loads use the same boot loader software. Lucent recommends that you always install a new boot loader when upgrading to a release.
<code>tnntrel.tar</code>	Tar file (T1 load) that contains images for the shelf controller and all MAX TNT slot cards.
<code>tnntrele.tar</code>	Tar file (E1 load) that contains images for the shelf controller and all MAX TNT slot cards.
<code>tnntbase.tar</code>	Tar file (T1 load) that contains all basic images required by a North American ISP and that is less than 8-MB in size. The file contains images for only the following modules: Shelf controller (T1), 8T1, UT1 (Frameline), T3, Ethernet-2, Ethernet-3, HDLC-2, 56K modem, and Series56 II modem.
<code>tnntbasee.tar</code>	Tar file (E1 load) that contains all basic images required by a European ISP and that is less than 8-MB in size. The file contains images for only the following modules: Shelf controller (E1), 8E1, UE1 (E1 Frameline), Ethernet-2, Ethernet-3, HDLC-2, 56K modem, and Series56 II modem.

You can obtain the files you need from the anonymous FTP server <ftp.ascend.com>. If you need technical assistance, see “Customer Service” on page 3.

Local access to the unit recommended

Whenever you install system software, Lucent recommends that you access the unit through the shelf controller serial or LAN port rather than a slot card port.

32-MB JEDEC DRAM card required for this release

For this release, the MAX TNT requires a 32-MB JEDEC DRAM card (model number TNT-SP-DRAM-32). New MAX TNT units now ship standard with the 32-MB DRAM card.

The 32-MB JEDEC DRAM card is *not* hot swappable. To install the card, you must turn off power to the MAX TNT, insert the card and then power on the MAX TNT. For additional information about the card, contact your service representative.

Flash size limitations for this upgrade

Because the MAX TNT supports many different slot card types, the tar files containing slot-card code images can be too large to load on an 8-MB flash card. TAOS 7.0.0 introduced a new Load-Select profile type that prevents loading the entire set of slot-card images. The profile causes the system to determine which card types are present and load only those images. For details about the Load-Select profile, see the *MAX TNT Reference Guide*.

In addition, in this release, the `tntbase.tar` and `tntbasee.tar` files are less than 8-MB in size and are guaranteed to fit on an 8-MB flash card.

If neither of the small tar files are appropriate for your systems, to load this release to 8MB flash, make sure that all parameters in the Load-Select profile are set to `auto` and that the combined binaries required to run the system and its cards do not exceed 8MB. Following are the approximate sizes of each binary in the tar file:

Table 14. Approximate sizes of shelf controller and card binaries

System component	Binary filename	Approx. size (KB)
Shelf controller (T1)	<code>tntsr/tntsr.ffs</code>	1800
Shelf controller (E1)	<code>tntsre/tntsre.ffs</code>	1800
8T1	<code>tnt8t1/tnt8t1.ffs</code>	275
UTI (Frameline)	<code>tntut1/tntut1.ffs</code>	825
8E1	<code>tnt8e1/tnt8e1.ffs</code>	260
UE1 (E1 Frameline)	<code>tntue1/tntue1.ffs</code>	810
T3	<code>tntt3/tntt3.ffs</code>	310
Ethernet-2	<code>tntenet2/tntenet2.ffs</code>	240
Ethernet-3	<code>tntenet3/tntenet3.ffs</code>	355
HDLC-2	<code>tnthdlc2/tnthdlc2.ffs</code>	1005
HDLC-2EC	<code>tnthdlc2ec/tnthdlc2ec.ffs</code>	1000
SWAN	<code>tntswan/tntswan.ffs</code>	725
UDS3	<code>tntuds3/tntuds3.ffs</code>	730
DS3-ATM	<code>tntds3atm/tntds3atm.ffs</code>	735
OC3-ATM	<code>tntoc3atm/tntoc3atm.ffs</code>	730
Analog modem	<code>tntamdm/tntamdm.ffs</code>	700
56K modem	<code>tntmdm56k/tntmdm56k.ffs</code>	850
Series56 I/ Series56 II	<code>tntcsmx/tntcsmx.ffs</code>	990
Series56 III	<code>tntcsmv/tntcsmv.ffs</code>	980

Table 14. Approximate sizes of shelf controller and card binaries

System component	Binary filename	Approx. size (KB)
MultiDSP	tntmadd/tntmadd.ffi	1300
STM-0	tntstm0/tntstm0.ffi	300

Saving the system configuration

As a general practice, always save the system configuration before upgrading or downgrading system software. You can then restore the configuration along with earlier system software if anything unexpected occurs during the upgrade or downgrade. If you use TFTP to save the system configuration, the target file must exist on the TFTP server and you must have permission to write it. For example, the following commands executed on a TFTP server create a target file and set its permissions:

```
$ touch /tftpboot/config/testcfg.1
$ chmod a=rw /tftpboot/config/testcfg.1
```

Before you save the system configuration, you must enable the Allow-Password permission in the MAX TNT User profile to save the configured passwords. If you do not have Allow-Password permission enabled, you will be prompted to confirm that you wish to save the configuration without passwords. If you do so and then restore the saved configuration, all passwords in the configuration are wiped out. The following commands executed on the MAX TNT save the system's configuration to the target file on the TFTP server and then restore the saved configuration:

```
admin> save -a network 10.10.10.10 /tftpboot/config/testcfg.1
admin> load config network 10.10.10.10 /tftpboot/config/testcfg.1
```

Upgrade instructions

These instructions show how to upgrade to MAX TNT TAOS 8.0.0 from TAOS version 7.0.0 or later. If you are not sure which version the system is running, enter the `version` command. For example:

```
admin> version
Software version 7.2.0
```

If the system is running a software version earlier than 7.0.0, please upgrade to 7.0.0 first and then follow the instructions in this note. For information about upgrading to 7.0.0, you can access the MAX TNT TAOS 7.0.0 release note at <http://www.ascend.com/doclibrary>.

Note: Under certain conditions, the `load tar` command might recognize no slot cards and load only the shelf controller image during the upgrade procedure. If this occurs, reset the system and load the tar file again. The second `load tar` command will load the appropriate slot-card images for the system.

Upgrading a standalone MAX TNT unit

To upgrade a standalone unit with 8MB flash, proceed as follows:

- 1 Log into the system and save its configuration to a TFTP server. This step is optional but strongly recommended. For details, see "Saving the system configuration" on page 25.

Upgrade and downgrade procedures

Upgrade instructions

- 2 Verify that the combined binaries required to run the system and its cards do not exceed 8MB. See “Approximate sizes of shelf controller and card binaries” on page 24.
- 3 Verify that the Load-Select profile is configured to automatically load only required binaries. All parameters in the profile must be set to `auto`.
- 4 Format the flash card. For example:

```
admin> format flash-card-1
```
- 5 Load the boot loader. For example:

```
admin> load boot-sr network 10.10.10.10 tntsrbin
```
- 6 Load the tar file. For example:

```
admin> load tar network 10.10.10.10 tntrel.tar
```
- 7 Reset the system. This step is required. For example:

```
admin> reset
```
- 8 Telnet into the system via the serial connection. Verify that the shelf controller IP address is set. For example:

```
admin> get ip-interface { { 1 c 1 } 0 } ip-address
[in IP-INTERFACE/{ { shelf-1 controller 1 } 0 } :ip-address]
ip-address = 10.10.10.2/24
```

If the address is not set, open the IP-Interface profile for the shelf controller and set the address. For example:

```
admin> read ip-interface { { 1 c 1 } 0 }
IP-INTERFACE/{ { shelf-1 controller 1 } 0 } read

admin> set ip-address = 10.10.10.2/24

admin> write
IP-INTERFACE/{ { shelf-1 controller 1 } 0 } written
```
- 9 Load the system configuration. This step is optional, but recommended. For example:

```
admin> load config network 10.10.10.10 /tftpboot/config/tntconfig
```
- 10 Format the flash card again. For example:

```
admin> format flash-card-1
```
- 11 Load the tar file again. For example:

```
admin> load tar network 10.10.10.10 tntrel.tar
```
- 12 Reset the system. This step is optional, but recommended. For example:

```
admin> reset
```

Upgrading a multishelf MAX TNT unit

If you are upgrading a multishelf system, you must propagate the new boot loader to the slave shelves by using the Loadslave command. (The version of the `tntsrbin` file on the master shelf must match the `tntsrbin` version on the slave shelves. Otherwise, the slave shelves cannot load code from the master shelf.) In addition, you must propagate a link to a redundant image of the tar file located in onboard flash.

To upgrade a multishelf unit with 8MB flash, proceed as follows:

- 1 Log into the master shelf and save the configuration to a TFTP server. This step is optional but strongly recommended. For details, see “Saving the system configuration” on page 25.

- 2 Verify that the combined binaries required to run the system and its cards do not exceed 8MB. See “Approximate sizes of shelf controller and card binaries” on page 24.
- 3 Verify that the Load-Select profile is configured to automatically load only required binaries. All parameters in the profile must be set to `auto`.
- 4 Format the flash card. For example:

```
admin> format flash-card-1
```
- 5 Load the boot loader. For example:

```
admin> load boot-sr network 10.10.10.10 tntsrbin
```
- 6 Propagate the new boot loader to the slave shelves. For example, the following command propagates the boot loader to a slave shelf with a rotary-switch setting of 2:

```
admin> loadslave 2 boot-sr
```
- 7 Load the tar file. For example:

```
admin> load tar network 10.10.10.10 tntrel.tar
```
- 8 Reset the system. This step is required. For example:

```
admin> reset -a
```
- 9 Telnet into the system via the serial connection. Verify that the master shelf controller IP address is set. For example:

```
admin> get ip-interface { { 1 c 1 } 0 } ip-address  
[in IP-INTERFACE/{ { shelf-1 controller 1 } 0 } :ip-address]  
ip-address = 10.10.10.2/24
```

If the address is not set, open the IP-Interface profile for the shelf controller and set the address. For example:

```
admin> read ip-interface { { 1 c 1 } 0 }  
IP-INTERFACE/{ { shelf-1 controller 1 } 0 } read  
  
admin> set ip-address = 10.10.10.2/24  
  
admin> write  
IP-INTERFACE/{ { shelf-1 controller 1 } 0 } written
```
- 10 Load the system configuration. This step is optional, but recommended. For example:

```
admin> load config network 10.10.10.10 /tftpboot/config/tntconfig
```
- 11 Format the flash card again. For example:

```
admin> format flash-card-1
```
- 12 Load the tar file again. For example:

```
admin> load tar network 10.10.10.10 tntrel.tar
```
- 13 Use the Loadslave command to propagate a link to the `image2` file, which is a redundant image of the tar file created in onboard flash. For example, the following command propagates the image to a slave shelf with a rotary-switch setting of 2:

```
admin> loadslave 2 image2
```
- 14 Reset the system. This step is optional, but recommended. For example:

```
admin> reset -a
```

Downgrade instructions

Because releases are not necessarily backward compatible, Lucent recommends that you always restore a backup configuration made under the previous version or one of its predecessors.

If you have enabled extended profiling and then must downgrade to an earlier software version, see “Additional onboard memory for extended profiling” on page 90, for important information.

Note: Serial access to the MAX TNT unit is required for downgrading to a previous release from MAX TNT TAOS 8.0.0. Because of the new profiles and functionality introduced in MAX TNT TAOS 8.0.0, you must initialize the system by clearing the onboard nonvolatile random access memory (NVRAM) when performing a downgrade. When you clear NVRAM, the initialized system starts up unconfigured, just as it was when you first installed it, with no IP address assignments.

Downgrading a standalone MAX TNT unit

To restore an earlier system software version, proceed as follows:

- 1 Log into the MAX TNT and save the current configuration to a TFTP server. This step is optional, but recommended.
- 2 Reformat the flash card to the old format. This is required. For example:

```
admin> format -o flash-card-1
```
- 3 Load the previous version of the boot loader. For example:

```
admin> load boot-sr network 10.10.10.10 tntsrb.bin
```
- 4 Load the previous version of the tar file. For example, to load via TFTP from a local host:

```
admin> load tar network 10.10.10.10 tntrel.tar
```
- 5 Clear NVRAM. This step is required when downgrading. For example:

```
admin> nvram -f
```
- 6 Telnet into the system via the serial connection. Open the IP-Interface profile for the shelf controller and set the address. For example:

```
admin> read ip-interface { { 1 c 1 } 0 }  
IP-INTERFACE/{ { shelf-1 controller 1 } 0 } read  
admin> set ip-address = 10.10.10.2/24  
admin> write  
IP-INTERFACE/{ { shelf-1 controller 1 } 0 } written
```
- 7 Load a backup configuration made under the restored software version or one of its predecessors. For example:

```
admin> load config network 10.10.10.10 /tftpboot/config/7x-config
```
- 8 Reset the system. This step is optional, but recommended. For example:

```
admin> reset
```

Downgrading a multishelf MAX TNT unit

If you are downgrading a multishelf system, you must propagate the restored boot loader to the slave shelves by using the Loadslave command. (The version of the `tntsr.b` file on the master shelf must match the `tntsr.b` version on the slave shelves. Otherwise, the slave shelves cannot load code from the master shelf.) In addition, you must propagate a link to a redundant image of the restored tar file. To restore an earlier system software version, proceed as follows:

- 1 Log into the master shelf and save the current configuration to a TFTP server. This step is optional, but recommended.
- 2 Reformat the flash card to the old format. This is required. For example:

```
admin> format -o flash-card-1
```
- 3 Load the previous version of the boot loader. For example:

```
admin> load boot-sr network 10.10.10.10 tntsr.b
```
- 4 Propagate the boot loader to the slave shelves. For example, the following command propagates the boot loader to a slave shelf with a rotary-switch setting of 2:

```
admin> loadslave 2 boot-sr
```
- 5 Load the previous version of the tar file. For example, to load via TFTP from a local host:

```
admin> load tar network 10.10.10.10 tntrel.tar
```
- 6 Clear NVRAM. This step is required when downgrading. For example:

```
admin> nvram -f
```
- 7 Telnet into the system via the serial connection. Open the IP-Interface profile for the shelf controller and set the address. For example:

```
admin> read ip-interface { { 1 c 1 } 0 }  
IP-INTERFACE/{ { shelf-1 controller 1 } 0 } read  
admin> set ip-address = 10.10.10.2/24  
admin> write  
IP-INTERFACE/{ { shelf-1 controller 1 } 0 } written
```
- 8 Load a backup configuration made under the restored software version or one of its predecessors. For example:

```
admin> load config network 10.10.10.10 /tftpboot/config/7x-config
```
- 9 Use the Loadslave command to propagate a link to the `image2` file, which is a redundant image of the tar file created in onboard flash. For example, the following command propagates the image to a slave shelf with a rotary-switch setting of 2:

```
admin> loadslave 2 image2
```
- 10 Reset the system. This step is optional, but recommended. For example:

```
admin> reset -a
```

Upgrade and downgrade procedures

Downgrade instructions

Built-in features in MAX TNT TAOS 8.0.0

Modem Manager

Firmware versions for digital modems

The Conexant firmware versions for MAX TNT Digital Modem cards include support for V.90, K56flex, K56plus, and all slower, standard modem speeds. This release includes the following Conexant firmware:

- Series56 Digital Modem cards (also called CSM/1, TNT-SL-48MOD-S56) support V2.0982-K56_2M_DLP_CSM firmware.
- Series56 II Digital Modem cards (also called CSM/3, TNT-SL-48MOD-SGL and TNT-SL-48MOD-S-C) support V5.817 firmware.
- Series56 III Digital Modem cards (also called CSM/3V, TNT-SL-48MODV3-S-C) support V5.8173 firmware.

The V5.817 and V5.8173 firmware include a fix for synchronization rate failures with some PCtel chipset modems. The V5.8173 firmware also provides a fix for synchronization failures observed with some Lucent winmodems.

Firmware versions for MultiDSP cards

This release includes the following Lucent firmware versions for MultiDSP cards:

- 48-port MultiDSP cards (TNTP-SL-ADI-C or TNTV-SL-ADI-C) support Lucent V1.16.14 firmware.
- 96-port MultiDSP cards (APX8-SL-96DSP) support Lucent V1.16.14 firmware.

Series56 III modem card support

The Series56 III Digital Modem card (TNT-SL-48MODV3-S-C) is a single-slot 48-port card that is the functional equivalent of the Series56 II card. Ongoing support continues in parallel for the Series56, Series56 II, and Series56 III modules.

The new Series56 III has the same installation and configuration procedures as the Series56 II card, described in the *MAX TNT Hardware Installation Guide*. The procedures are also described in the Series56 II guide, which you can access online after registering at <http://www.ascend.com/doclibrary>.

The output of the Show command identifies the Series56 III Digital Modem card as csmv-card, as shown in the following example:

```
admin> show
Shelf 1 ( standalone ):
    { shelf-1 slot-14 0 }      UP      csmv-card
```

Expanded MultiDSP card support

In addition to the 48-port MultiDSP card (TNT-P-SL-ADI-C or TNTV-SL-ADI-C) a 96-port MultiDSP card (APX8-SL-96DSP) is now available.

Modem service is now supported and enabled by default on both MultiDSP cards.

With the appropriate software licenses, services currently supported on the 48-port MultiDSP card are: modem (for example, V.90), ISDN (HDL), V.110, PHS, and VoIP (voice). The 96-port card does not support PHS or VoIP in this release.

PHS functionality now supports a fixed data rate of 32Kbps (PIAFS 1.0), or a fixed data rate of either 32Kbps or 64Kbps for the duration of a call (PIAFS 2.0), or a data rate that switches between 32Kbps and 64Kbps during a call, depending on what the wireless bandwidth permits (PIAFS 2.1). The PIAFS 2.1 functionality requires a separate license (see “PHS Internet Access Forum Standard (PIAFS) 2.1 on MultiDSP” on page 175).

The 48-port MultiDSP card supports 48 ports of any service and handles up to two different services per card. In this release, when running two services per card, the services can be used only in one of the following combinations:

- Data (modem/ISDN) with V.110
- Data (modem/ISDN) with PHS
- Data (modem/ISDN) with VoIP

The 96-port MultiDSP card currently supports 96 ports of data (modem/ISDN) and/or V.110 service, and handles up to two different services per card. When running two services per card, one service must be data and the other must be V.110. The 96-port card does not support PHS or VoIP in this release.

In this release, the following configuration restrictions apply:

- The 96-port and 48-port MultiDSP card cannot be used together in the same unit.
- The dual-port Series56 card (TNT-SL-48MOD-S56) cannot be used in the same unit with MultiDSP cards.

Multiple 48-port MultiDSP cards can be used in the same unit, and the Series56 II (TNT-SL-48MOD-SGL and TNT-SL-48MOD-S-C) and Series56 III (TNT-SL-48MODV3-S-C) cards can be used in the same unit as a MultiDSP card.

For further details on the MultiDSP cards, see the MultiDSP guide at <http://www.ascend.com/doclibrary>. After you register, you can view or download the guide.

Stac compression for asynchronous calls

Stac and Microsoft/Stac (MS-Stac) compression are now supported for asynchronous Point-to-Point Protocol (PPP) connections received on the Series56 II, Series56 III, and MultiDSP modems.

Note: This feature is *not* supported on the Series56 dual-slot card.

If a caller requests Stac or MS-Stac compression and the caller’s profile is configured properly, the requested compression method is negotiated for the modem connection. The following types of PPP link compression are supported:

- Stac compression uses a modified version of draft 0 of the Compression Control Protocol (CCP), which predates RFC 1974. Older Ascend equipment supports this compression method. It is not recommended for use with IPX connections.
- Stac-9 compression uses draft 9 of the Stac LZS compression protocol, which is described in RFC 1974. Most devices use this compression method.
- MS-Stac compression is the method used by Microsoft Windows 95 clients.

For details about configuring PPP connections, modem connections, and link compression, see the *MAX TNT Network Configuration Guide*.

Settings in Connection profiles

Following is the parameter, shown with its default setting, for specifying PPP link compression in a local Connection profile:

```
[in CONNECTION/" ":ppp-options]
link-compression = stac
```

Parameter	Specifies
Link-Compression	Link-compression method to use for PPP-encapsulated packets transmitted and received on the connection. During the negotiation phase of the connection, both sides must agree to use the specified method. Supported values are <code>stac</code> (the default), <code>stac-9</code> , and <code>ms-stac</code> .

For example, the following commands configure MS-Stac compression:

```
admin> read connection windows-user
CONNECTION/windows-user read

admin> set ppp-options link-compression = ms-stac

admin> write
CONNECTION/windows-user written
```

Settings in RADIUS user profiles

RADIUS uses the following attribute-value pair for setting link compression:

RADIUS attribute	Value
Ascend-Link-Compression (233)	Link-compression method to use for PPP-encapsulated packets transmitted and received on the connection. During the negotiation phase of the connection, both sides must agree to use the specified method. The following values are supported: Link-Comp-Stac (1). (Stac compression.) Link-Comp-Stac-Draft-9 (2). (Stac-9 compression.) Link-Comp-MS-Stac (3). (MS-Stac compression.)

Following is a sample RADIUS user profile that uses Stac-9 compression:

```
user-1 Password = "localpw"
      Service-Type =Framed-User,
      Framed-Protocol = PPP,
```

```
Ascend-Link-Compression = Link-Comp-Stac-Draft-9,  
Framed-IP-Address = 10.1.1.1,  
Framed-IP-Netmask = 255.255.255.0
```

Note: When the Use-Answer-For-All-Defaults parameter in the MAX TNT Answer-Defaults profile is set to *yes* (the default), attributes that are not specified in callers' RADIUS profiles take their value from Answer-Defaults settings. You can set the following parameter, shown with its default value, to specify the link compression type for all incoming PPP calls:

```
[in ANSWER-DEFAULTS:ppp-answer]  
link-compression = none
```

If you set this parameter to a compression method (such as *stac-9*), the MAX TNT negotiates only that type of compression for incoming calls. However, if a caller requests another compression method (such as *MS-Stac*), the MAX TNT establishes the link but does not use compression. For example, the following commands cause the system to negotiate only Stac-9 compression for all incoming PPP calls:

```
admin> read answer-defaults  
ANSWER-DEFAULTS read  
  
admin> set ppp-answer link-compression = stac-9  
  
admin> write  
ANSWER-DEFAULTS written
```

WAN Access Server

Ethernet-3 slot card support

The Ethernet-3 card (TNT-SL-E100-V-C) was introduced with limited availability in earlier 7.x releases. It is now generally available in MAX TNT TAOS 8.0.0.

The Ethernet-3 card is a high performance Ethernet module with one 100-MB interface designed for demanding applications such as VoIP. The Ethernet-3 card uses the same installation procedures as the Ethernet-2 card (TNT-SL-E10-100), which is described in the *MAX TNT Hardware Installation Guide*. For details about using the card with VoIP, see "Ethernet requirements for VoIP processing" on page 256.

The output of the Show command identifies the Ethernet-3 card as the *ether3* card, as shown in the following example:

```
admin> show  
Shelf 1 ( standalone ):  
  { shelf-1 slot-14 0 }      UP      ether3
```

OC3-ATM slot card support

The OC3-ATM card was introduced in earlier 7.x releases. Support for ATM SVCs is new in this release.

The MAX TNT OC3 (Copper) ATM slot card (TNT-SL-OC3-C) and MAX TNT OC3 (Fibre) ATM slot card (TNT-SL-OC3-F) support the following features:

- One Optical Carrier 3 (OC-3) unchannelized port
- IP routing over ATM

- RFC 1483 (multiprotocol encapsulation over ATM)
- Layer 2 permanent virtual circuit (PVC) switching between ATM and Frame Relay
- Converting from multiprotocol encapsulation over ATM (RFC 1483) to multiprotocol encapsulation over Frame Relay (RFC 2427)
- *FRF.8 Frame Relay ATM/PVC Service Interworking Implementation Agreement*
- Operation, Administration, and Maintenance (OAM) F4/F5
- ATM switched virtual circuits (SVCs)

To access the OC3-ATM configuration guide, go to <http://www.ascend.com/doclibrary>. After you register, you can view or download the guide.

HDLC2-EC-C slot card support

The Hybrid Access III card (TNT-SL-HDLC2-EC-C) is similar to the Hybrid Access II card (TNT-SL-HDLC2), but the new card includes hardware support for encryption and compression. Ongoing support continues in parallel for both cards.

Both Hybrid Access cards provide High-level Data Link Control (HDLC) processing for inbound digital calls and support up to 186 HDLC sessions. The new HDLC2-EC-C card uses the same installation procedures as the HDLC2 card, which is described in the *MAX TNT Hardware Installation Guide*.

The output of the Show command identifies the HDLC2-EC-C card as `hdlc2ec` card, as shown in the following example:

```
admin> show
Shelf 1 ( standalone ):
  { shelf-1 slot-14 0 }      UP      hdlc2ec
```

STM-0 slot card support

Support for the Synchronous Transport Module (STM)-0 card (TNT-SL-STM0) is new in MAX TNT TAOS 8.0.0. The card is an optical 51.85 Mbps communication circuit designed to be used with an SS7 signaling gateway.

All 28 T1 lines can be configured as SS7 data trunks. When configured as an SS7 data trunk, the ASG takes control of the data trunks, telling the MAX TNT when to bring calls up or down. Note that the STM card does not support Call-Routing profiles, ISDN Primary Rate Interface (PRI) signaling, or inband signaling.

To access the STM-0 configuration guide, go to <http://www.ascend.com/doclibrary>. After you register, you can view or download the guide.

Bandwidth Allocation Control Protocol (BACP)

MAX TNT units now support BACP for PPP Multilink Protocol (MP) connections. MP is described in RFC 1990. In previous releases, MP connections could use multiple channels, but the bandwidth allocation was static.

BACP is described in RFC 2125. It provides dynamic bandwidth allocation based on a utilization threshold, using criteria that are very similar to those used by the bandwidth-on-

demand feature in Multilink Protocol Plus (MP+). BACP can be used with digital or analog links.

For dynamic bandwidth allocation to work on an MP connection, both sides of the connection must support BACP. The following parameters (shown with sample settings) enable BACP:

```
[in ANSWER-DEFAULTS:mp-answer]
bacp-enable = yes

[in CONNECTION/"":mp-options]
bacp-enable = yes
```

Parameter	Specifies
BACP-Enable	Enable/disable BACP for MP connections. BACP is disabled by default. In the Answer-Defaults profile, the <code>yes</code> setting enables the system to accept an MP call that requests BACP bandwidth management. In a Connection profile, it enables a specific connection to use BACP bandwidth management.

BACP shares the parameters used by MP+ to specify criteria for adding or subtracting bandwidth. Following are the relevant parameters, shown with default settings:

```
[in CONNECTION/"":mpp-options]
dynamic-algorithm = quadratic
bandwidth-monitor-direction = transmit
increment-channel-count = 1
decrement-channel-count = 1
seconds-history = 15
add-persistence = 5
sub-persistence = 10
target-utilization = 70
```

Information about how these parameters work together to affect bandwidth allocation is provided in the *MAX TNT Network Configuration Guide*. Details about each parameter are provided in the *MAX TNT Reference Guide*. Following is an example of configuring the system to enable BACP and of configuring an MP connection that uses BACP:

```
admin> read answer
ANSWER-DEFAULTS read

admin> set mp-answer bacp-enable = yes

admin> write
ANSWER-DEFAULTS written

admin> read CONNECTION mp-test
CONNECTION/mp-test read

admin> set encapsulation-protocol = mp

admin> set mp-options bacp-enable = yes

admin> set mp-options maximum-channels = 4

admin> set mpp-options bandwidth-monitor-direction = transmit-recv

admin> set mpp-options seconds-history = 30

admin> set mpp-options add-persistence = 10

admin> write
CONNECTION/mp-test written
```

X.75 frame size increase to 2048 bytes

For X.75 calls that use a Hybrid Access card, the supported frame size for data (excluding headers) is increased to 2048 bytes from 1532 bytes. Following are the relevant parameters, shown with their default values:

```
[in ANSWER-DEFAULTS:x75-answer]
frame-length = 1024

[in CONNECTION/"":x75-options]
frame-length = 1024
```

Parameter	Specifies
Frame-Length	Frame length to use for incoming X.75 connections. In previous releases, the maximum frame length was 1532 bytes. The valid range is now from 128 to 2048.

Support for multidigit trunk groups

The number of digits that can be specified in trunk groups is now configurable system-wide from 1 (the default) to 4. Following are MAX TNT parameters related to trunk group assignments, shown here with default values:

```
[in SYSTEM]
use-trunk-groups = no
num-digits-trunk-groups = 1

[in T1/{ any-shelf any-slot 0 }:line-interface:channel-config 1]
trunk-group = 9

[in E1/{ any-shelf any-slot 0 }:line-interface:channel-config 1]
trunk-group = 9

[in SWAN/{ any-shelf any-slot 0 }:line-config]
trunk-group = 9

[in CALL-ROUTE//{ { { any-shelf any-slot 0 } 0 } 0 }
trunk-group = 0
```

Parameter	Specifies
Use-Trunk-Groups	Enable/disable the use of trunk groups in the MAX TNT. If set to no (the default), the Num-Digits-Trunk-Groups and Trunk-Group settings do not apply. If set to yes, all channels must be assigned trunk group numbers.
Num-Digits-Trunk-Groups	Number of digits to allow for trunk groups, from 1 to 4. With the default of 1 (single-digit), trunk group numbers range from 2 to 9, and the dial-out telephone number is preceded by a single-digit number. For example, to dial the number 555-1212 on trunk 7, the dial-out telephone string is 75551212. If Num-Digits-Trunk-Groups is set to 2, 3, or 4, trunk group numbers range can include the specified number of digits (up to 9999), and the dial-out telephone number is always preceded by that number of digits. For example, if it is set to 2, to dial the number 555-1212 on trunk 7, the dial-out telephone string is 075551212.

Parameter	Specifies
Trunk-Group	Trunk group number. For a network line, it assigns channels to a trunk group. In a Call-Route profile, it specifies that calls received on that trunk group will be routed to shelf, slot, and port specified in the profile's index.

Note: When the MAX TNT is configured to interoperate with an external application for dial-out, the external system and the MAX TNT *must agree* about the number of digits in a trunk group number, or telephone numbers will not be parsed correctly and calls will fail.

Example of using 2-digit trunk groups

Following is an example that configures the MAX TNT to use 2-digit trunk groups and assigns the trunk group 9 to channels 1 through 5 of a T1 line:

```
admin> read system
SYSTEM read

admin> set num-digits-trunk-groups = 2

admin> write
SYSTEM read

admin> read t1 { 1 2 2 }
T1/{ shelf-1 slot-2 2 } read

admin> set line channel 1 trunk = 9
admin> set line channel 2 trunk = 9
admin> set line channel 3 trunk = 9
admin> set line channel 4 trunk = 9
admin> set line channel 5 trunk = 9

admin> write
T1/{ shelf-1 slot-2 2 } written

admin> read call-route { { { 1 2 0 } 0 } 0 }
CALL-ROUTE/{ { { shelf-1 slot-2 0 } 0 } 0 } read

admin> set trunk-group = 9

admin> write
CALL-ROUTE/{ { { shelf-1 slot-2 0 } 0 } 0 } written
```

Note about external dial-out applications

If an dial-out application uses trunk group assignments to select an available channel, it precedes each dial-out string with a trunk number. If both the application and the MAX TNT are configured for multidigit trunk groups, dial-out telephone numbers include the specified number of trunk number digits, up to a maximum of four. For example, if the systems are configured to use 2-digit trunk groups, the application sends the string 075551212 to the MAX TNT to dial 555-1212 on trunk 7.

Note: When the MAX TNT is configured to interoperate with an external application for dial-out, the external system and the MAX TNT *must agree* about the number of digits in a trunk group number, or telephone numbers will not be parsed correctly and calls will fail.

MAX TNT stacking


You can configure the MAX TNT to participate as a member of a group of local units that manage MP or MP+ bundles jointly (a stack). A stack can include both MAX TNT and MAX™ units, provided that the MAX units support Stacking Protocol Version 3. (Stacking Protocol Version 3 is incompatible with earlier versions.) Some network constraints apply when you are stacking MAX TNT and MAX units together.

In this release, the MAX TNT does not support token card authentication of stacked bundles. In addition, stacking is not supported for MAX TNT multiself units.

How stacking works

The initial bandwidth of a connection is established when user equipment dials into a MAX TNT or MAX unit. After the link has been authenticated, if more bandwidth is requested, the system that dialed the initial link can dial another link to add bandwidth. When stacking is enabled, the new link can be handled by any one of the stack peers. In Figure 1, the calling MAX unit dials three links, each of which is answered by a different stacked unit (*peer*).

Figure 1. Stacked bundle consisting of three links



The MAX TNT or MAX unit that establishes the initial link is the *bundle owner*. The bundle owner manages the connection's traffic across the stack. Stack peers forward inbound traffic from the WAN to the bundle owner, and the bundle owner receives outbound traffic destined for the far end and distributes it to bundled links. To balance the load among all available WAN channels, outbound data packets are assigned to bundled links on a rotating basis.

The stacking control interface and stacking data interface can use the same Ethernet segment, but most sites use separate segments for performance and management reasons.

Stacking control interface

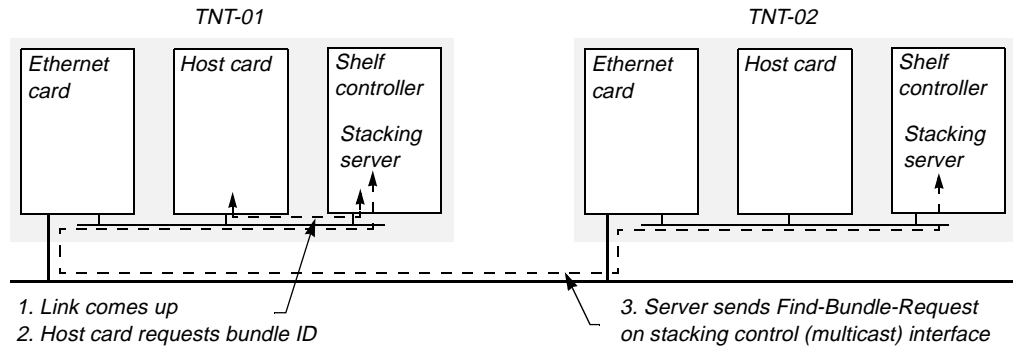
Stacking control packets are exchanged among stack peers to manage calls. When stacking is enabled and an MP or MP+ link has been authenticated (typically via RADIUS), stacking software on the MAX TNT host card requests a bundle ID from stacking software on the shelf controller (the stacking *server*).

If the server does not know where a bundle for the link resides (for example, if no bundle exists yet), it multicasts a Find-Bundle-Request packet on the Ethernet interface specified for stacking control packets. If the requesting server receives a reply that includes a bundle ID, it

adds the new link to that bundle. Otherwise, it assigns a bundle ID and becomes the bundle owner. At this point, the authentication acknowledgement is sent to the caller.

Figure 2 identifies the steps taken by the peer named TNT-01 to establish a bundle for an MP or MP+ link.

Figure 2. Stacking control messages to establish a bundle



Once a bundle has been established, additional stacking control messages are exchanged between peers. Some control packets are unicast to the system IP address of the target MAX TNT and the UDP port specified in the Stacking profile. Other stacking control packets are multicast on the Ethernet interface specified for stacking control packets (the multicast interface) to the specified multicast address, through the specified UDP port. Following are the relevant parameters, shown with their default settings:

```
[ STACKING/ " " ]
enabled = no
name = " "
udp-port = 5150
multicast-address = 239.192.74.72
multicast-interface-ip-address = 0.0.0.0
data-ip-address = 0.0.0.0
```


(For details, see “Stacking profile settings” on page 47.)

Note: You can configure any one of the system’s Ethernet interfaces as the stacking control interface. To avoid overloading the shelf controller, Lucent recommends choosing an interface other than the shelf controller Ethernet interface. All stack members must specify the same IP subnet as the multicast interface, and must agree on the multicast address and UDP port.

Stacking data interface

Once a link is established, stacking software on the MAX TNT host cards exchanges stacking data packets (user data) as well as keepalive packets (zero-length data packets). Figure 3 shows data and control packets being exchanged for a stacked session.

Figure 3. Stacking data and control packets



Depending on how the system address is set in the peers, unicast control packets can also be received on the stacking data interface. Following is the parameter for specifying the stacking data interface, shown with its default setting:

```
[ STACKING / " " ]  
data-ip-address = 0.0.0.0
```

(For details, see “Stacking profile settings” on page 47.) Stacking data packets are sent to the Data-IP-Address specified in the Stacking profile of the target MAX TNT and the UDP port that was negotiated between the stack peers at call setup time. The UDP port used by a host card is assigned when the card boots, and remains the same for all calls on the card.

Note: Stacking data packets are typically exchanged on an Ethernet interface other than the one used for stacking control packets.

Performance considerations

Performance testing indicates that, on a stack of eight MAX TNT units, calls were established without stacking-related timeouts at the rate of three to five calls per second per unit. The stack that was tested used the same Ethernet interface for stacking control and data packet. Approximately 60% of the calls were a mix of single-channel PPP and two-channel MP calls that were not stacked, and approximately 40% were two-channel MP+ calls that were stacked. Two RADIUS servers were used to authenticate the calls and process accounting records.

Many variables can affect stack performance, including the physical network configuration, percentage of calls that are multilink and therefore make use of the stack, speed of the RADIUS server(s), use of OSPF routing, and amount of user data transferred on each call.

If the shelf controller of a stacked MAX TNT becomes too busy, it starts delaying stacking control messages. The delay can prevent the system from responding to a Find-Bundle-Request for a bundle it owns. Failure to respond in time can cause the system that answered the call to form a new bundle rather than adding a link to an existing bundle. If this situation occurs, a message is sent to Syslog in the following format:

```
MMP-N: NetBundleReply too late failure. st=state
```

N is the MBID of the second bundle, and *state* is the state of the stacking session (usually *AwaitNcps*). This message indicates that the bundle owner replied too late, and that the second link will be torn down.


Network configuration considerations related to stacking

Typically, the stacking control interface can use a 10BaseT port, but the stacking data interface might require a dedicated full-duplex 100BaseT port for increased throughput. If the Ethernet performance is too slow, timeouts can occur when a system receives a large number of calls simultaneously. See “Performance considerations” on page 41 for related information.

Including an Ethernet switch for increased performance

If traffic is too high on one or both of the stacking interfaces, you can connect the affected segments in a star configuration with an Ethernet switch at the center. For example, Figure 4 shows a switch on the stacking data interface.

Figure 4. Stacking data interfaces in star configuration




Stacking MAX TNT and MAX units together

Because a MAX unit has only one Ethernet interface, you must use one of the following connection alternatives when stacking a MAX together with MAX TNT units:

- Use the same MAX TNT Ethernet interface for both stacking control and stacking data packets.
- Use a router to connect the interfaces, so the MAX can reach both the stacking control and stacking data interfaces.

Figure 5 shows a network configuration using a router between the control packet and data packet subnets:

Figure 5. Stacking MAX TNT and MAX units together



System IP address

The system IP address is an area of memory that contains the address of one of the Ethernet interfaces of the MAX TNT. By default, it contains the IP address assigned to the shelf controller Ethernet interface. If the System-IP-Addr has been set to the address assigned to a port on an Ethernet slot card, the system receives stacking control unicast on that interface.

(This behavior does not cause a problem for the stacked peers.) For more details, see the *MAX TNT Network Configuration Guide*.

Using the soft address for fault tolerance

The MAX TNT supports a soft IP interface, which is an internal interface that never goes down. Routing protocols always advertise the soft interface address as reachable on all interfaces that are up and running a routing protocol. The soft interface address is reachable as long as one of the system's interfaces is up and running a routing protocol.

Like the System-IP-Addr, the Data-IP-Address in the Stacking profile is an area of memory that contains the address of one of the Ethernet interfaces of the MAX TNT. If the specified interface becomes unavailable, all stacking data packets destined for the interface are lost.

Some applications use the soft interface to keep from being bound to a particular interface. To use the soft interface as the destination for stacking data packets, enter the soft IP interface address in the Data-IP-Address parameter in the Stacking profile.

Note: To receive data packets on any available Ethernet interface, you must configure multiple static routes in external routers, with each route specifying the soft address as the destination and the address of a physical Ethernet interface in the MAX TNT as the gateway address. (Make sure to set a high metric for the route through the shelf controller Ethernet interface, so that the interface is chosen only if all other Ethernet interfaces are down.)

For details about setting the soft address in the IP-Interface profile with the default index (which is reserved for that purpose), see the *MAX TNT Network Configuration Guide*.

Routing protocols and stacking

When a call arrives on a stacked peer and it is the first call of a bundle, routers learn the route to the caller through a routing protocol update. A subsequent call for the same bundle does not create a new route.

RIP and OSPF can both supply the required route information, but OSPF is faster. For this reason, Lucent recommends using OSPF when stacking is enabled. The performance difference between RIP and OSPF can be important when stacking is enabled. For example, suppose a caller dials into a stack and routes a packet out to the Internet. The instant a reply comes back from the Internet, the router must know which system has the call, so it can forward the reply packet to the proper system. If it forwards the reply packet to the wrong peer, that system will attempt to dial out a new call to the user.

Telco considerations related to stacking

If stack peers share a hunt group, incoming calls to the shared number can be answered by any one of the peers. Sharing a hunt group means that the same telephone number is specified in the channel configuration of the line profiles of each peer. For example, if a caller dials 555-1217 to access the site, the channel configurations might specify the 1217 or 17 telephone number, as shown in the following sample T1 profile. (The telephone number is the third value listed in each Channel-Config subprofile.)

```
admin> read t1 {1 2 1}
T1/{ shelf-1 slot-2 1 } read

admin> list line channel
[in T1/{ shelf-1 slot-2 1 }:line-interface:channel-config]
```

```
channel-config[1] = { switched-channel 9 17 { any-shelf any-slot 0+
channel-config[2] = { switched-channel 9 17 { any-shelf any-slot 0+
channel-config[3] = { switched-channel 9 17 { any-shelf any-slot 0+
channel-config[4] = { switched-channel 9 17 { any-shelf any-slot 0+
channel-config[5] = { switched-channel 9 17 { any-shelf any-slot 0+
channel-config[6] = { switched-channel 9 17 { any-shelf any-slot 0+
channel-config[7] = { switched-channel 9 17 { any-shelf any-slot 0+
channel-config[8] = { switched-channel 9 17 { any-shelf any-slot 0+
channel-config[9] = { switched-channel 9 17 { any-shelf any-slot 0+
channel-config[10] = { switched-channel 9 17 { any-shelf any-slot +
channel-config[11] = { switched-channel 9 17 { any-shelf any-slot +
channel-config[12] = { switched-channel 9 17 { any-shelf any-slot +
channel-config[13] = { switched-channel 9 17 { any-shelf any-slot +
channel-config[14] = { switched-channel 9 17 { any-shelf any-slot +
channel-config[15] = { switched-channel 9 17 { any-shelf any-slot +
channel-config[16] = { switched-channel 9 17 { any-shelf any-slot +
channel-config[17] = { switched-channel 9 17 { any-shelf any-slot +
channel-config[18] = { switched-channel 9 17 { any-shelf any-slot +
channel-config[19] = { switched-channel 9 17 { any-shelf any-slot +
channel-config[20] = { switched-channel 9 17 { any-shelf any-slot +
channel-config[21] = { switched-channel 9 17 { any-shelf any-slot +
channel-config[22] = { switched-channel 9 17 { any-shelf any-slot +
channel-config[23] = { switched-channel 9 17 { any-shelf any-slot +
channel-config[24] = { switched-channel 9 17 { any-shelf any-slot +
```

If stack peers have the same hunt group specification, any one of the units can answer an incoming call to that number. For details about how to assign hunt groups, see the *MAX TNT Hardware Installation Guide*.

User equipment that supports Multilink PPP Plus (MP+) or Bandwidth Allocation Control Protocol (BACP) can request a telephone number before dialing to add bandwidth. This capability enables the bundle owner to return the telephone number of one of its own available channels, to avoid use of the stack. However, if no WAN channels are free in the bundle owner, the system requests a telephone number from the stack. If it does not receive a reply within 0.5 seconds, the system sends the request again, up to a total of three requests. If it receives a telephone number, the system returns the number to the caller, which uses that number to dial the next link.

The fact that MP+ and BACP devices request a telephone number before dialing can be useful for keeping links together on a single system, even when hunt groups are shared across the stack. For example, after the first call has been answered, the device can request a private number (a number that is not part of the shared hunt group) to use for the second call, to ensure that the links are managed on one system.

Conversely, an MP device that does not support BACP uses the same telephone number for all links required to establish the call at its configured bandwidth. As a result, each link can be handled by a different system in the hunt group.


Authentication and accounting considerations related to stacking

RADIUS authentication is recommended when stacking is in use, because all of the stacked peers must have access to the same authentication information for incoming calls. If you cannot use RADIUS authentication, you must replicate the Connection profiles for multilink PPP calls on each of the stacked systems.

RADIUS authentication

In the configuration shown in Figure 6, if the initial link is answered by the peer labeled TNT-03, it authenticates the caller via RADIUS. If the caller requests additional bandwidth and the subsequent call is answered by TNT-02, it authenticates the caller by accessing the same profile on the RADIUS server.

Figure 6. Stacked peers authenticating calls via RADIUS



When the MAX TNT is using RADIUS to authenticate calls, all of the stacked peers must have access to the same RADIUS database information. (Routes and other pseudo-user profiles must not be system-specific.)

For details about configuring the MAX TNT to use RADIUS authentication, see the *MAX TNT RADIUS Guide*. Following is a sample RADIUS profile for a multilink call that can use up to eight channels:

```
max-01 Password = "localpw"
    Service-Type = Framed-User,
    Framed-Protocol = MPP,
    Framed-IP-Address = 10.10.10.64,
    Framed-IP-Netmask = 255.255.255.0,
    Ascend-Base-Channel-Count = 2,
    Ascend-Maximum-Channels = 8,
    Ascend-DBA-Monitor = DBA-Transmit-Recv,
    Ascend-Seconds-Of-History = 30,
    Ascend-Add-Seconds = 10,
    Ascend-Dial-Number = 555-1217
```

RADIUS accounting

For information about configuring the MAX TNT to use RADIUS accounting, see the *MAX TNT RADIUS Guide*. In RADIUS accounting Start and Stop records, the value of Ascend-Multilink-ID is the same for all channels of a connection, including stacked channels. As a result, all calls belonging to the same bundle can be identified by this value.

The Ascend-Num-In-Multilink value tracks the number of channels currently connected, even if those channels are spread across multiple machines in a stacked environment. This attribute appears in both Start and Stop records, starting with a value of 1 (one).

Note: In earlier releases, records sent by the MAX TNT reported an initial value of zero for the Ascend-Num-In-Multilink attribute. The value was incremented for each subsequent link and then decremented as appropriate in Stop records. These values were inconsistent with those sent to RADIUS accounting servers by other TAOS units, which assign an initial value of one.

The Ascend-Disconnect-Cause and Acct-Multi-Session-Id attributes are transferred properly between stacked sessions.

The Acct-Link-Count value tracks the highest number of channels connected, even if those channels are spread across multiple machines in a stacked environment.

The Framed-IP-Address shows the address of the caller if the accounting record belongs to the bundle owner. If the accounting record belongs to a stack peer, the IP address is 0.0.0.0.

Following are example records for the initial link of a stacked bundle:

```
Mon Dec 6 15:45:52 1999
  User-Name = "max-01"
  NAS-Identifier = 10.1.1.2
  NAS-Port = 18
  NAS-Port-Type = Sync
  Acct-Status-Type = Start
  Acct-Delay-Time = 0
  Acct-Session-Id = "286464519"
  Acct-Authentic = RADIUS
  Ascend-Multilink-ID = 276234242
  Ascend-Num-In-Multilink = 1
  Acct-Link-Count = 0
  Acct-Multi-Session-Id = "10770002"
  Ascend-Modem-PortNo = 2
  Ascend-Modem-SlotNo = 2
  Ascend-Modem-ShelfNo = 1
  Caller-Id = "1119855029"
  Client-Port-DNIS = "3805"
  Framed-Protocol = MPP
  Framed-IP-Address = 10.10.10.64
```

```
Mon Dec 6 15:46:42 1999
  User-Name = "max-01"
  NAS-Identifier = 10.1.1.2
  NAS-Port = 18
  NAS-Port-Type = Sync
  Acct-Status-Type = Stop
  Acct-Delay-Time = 0
  Acct-Session-Id = "286464519"
  Acct-Authentic = RADIUS
  Acct-Session-Time = 50
  Acct-Input-Octets = 35007
  Acct-Output-Octets = 35022
  Acct-Input-Packets = 50
  Acct-Output-Packets = 52
  Ascend-Disconnect-Cause = 185
  Ascend-Connect-Progress = 60
  Ascend-Xmit-Rate = 56000
  Ascend-Data-Rate = 56000
  Ascend-PreSession-Time = 4
  Ascend-Pre-Input-Octets = 237
  Ascend-Pre-Output-Octets = 200
  Ascend-Pre-Input-Packets = 8
  Ascend-Pre-Output-Packets = 8
  Ascend-First-Dest = 10.1.1.2
  Ascend-Multilink-ID = 276234242
  Ascend-Num-In-Multilink = 0
  Acct-Link-Count = 2
  Acct-Multi-Session-Id = "10770002"
```

```
Ascend-Modem-PortNo = 2
Ascend-Modem-SlotNo = 2
Ascend-Modem-ShelfNo = 1
Caller-Id = "1119855029"
Client-Port-DNIS = "3805"
Framed-Protocol = MPP
Framed-IP-Address = 10.10.10.64
```

Stacking profile settings

Following are the parameters, shown with default values, for configuring stacking:

```
[ STACKING/" " ]
enabled = no
name = " "
udp-port = 5150
multicast-address = 239.192.74.72
multicast-interface-ip-address = 0.0.0.0
data-ip-address = 0.0.0.0
```

Parameter	Specifies
Enabled	Enable/disable stacking. If set to <code>yes</code> , stacking is enabled. If set to <code>no</code> (the default), none of the other settings in the Stacking profile apply. To disable a stack, set Enabled to No for each of the stack peers.
Name	Text string, up to 16 characters, which names the stack. All members of a stack specify the same name. Stacking control packets include this string to identify members of the same stack. Multiple stacks can exist on the same Ethernet segment if the stacks have different names.
UDP-Port	UDP port number to use for intrastack control packets. The default is 5150. All members of a stack must specify the same UDP port number. Multiple stacks can specify the same port number because the port does not have to be unique to a stack.
Multicast-Address	Multicast destination address for multicast stacking control packets. The packets are sent to the specified multicast address and UDP port number (immediately above). The default setting is 239.192.74.72, which is within the organization local scope defined in RFC 2365 as the address space from which an organization should allocate subranges when defining scopes for private use. The specified address must be a valid multicast (class D) address.
Multicast-Interface-IP-Address	IP address of the Ethernet port to be used for stacking IP multicast control traffic. With the default zero address, the system's shelf controller Ethernet interface is used.
Data-IP-Address	IP address of the Ethernet port to be used for stacking data traffic. The system advertises the address to other members of the stack in stacking control packets, and those systems in turn send stacking data packets to that address. With the default zero address, the System-IP-Addr is advertised instead. Enter the soft IP interface address for fault tolerance. (See "Using the soft address for fault tolerance" on page 43.)


You can add a system to an existing stack by configuring its Stacking profile and other setups. You need not reboot the system, and stack operation is not affected. Because a stack is a collection of peers, none keeps a list of the stack membership.

You can remove a system from a stack by disabling stacking in its Stacking profile. Existing stacked calls continue to function normally, but the system cannot operate as a stack peer for any new calls.

Example of stacking MAX TNT units

The following example shows how to stack two MAX TNT units. The stacking control interface is a 10BaseT Ethernet segment, and the full-duplex ports of the Ethernet cards in both units are connected to a 100BaseT network, as shown in Figure 7.

Figure 7. Example stack of two MAX TNT units



This example shows sample configurations for the two-system stack shown in Figure 7. Both systems receive calls on T1 or T3 lines. The system labeled TNT-01 supports the following slot cards:

```
admin> show
Shelf 1 ( standalone ):
  { shelf-1 slot-1 0 }      UP      t3-card
  { shelf-1 slot-2 0 }      UP      t3-card
  { shelf-1 slot-3 0 }      UP      4ether2-card
  { shelf-1 slot-4 0 }      UP      4ether2-card
  { shelf-1 slot-5 0 }      UP      hdlc2-card
  { shelf-1 slot-6 0 }      UP      hdlc2-card
  { shelf-1 slot-7 0 }      UP      hdlc2-card
  { shelf-1 slot-8 0 }      UP      hdlc2-card
  { shelf-1 slot-12 0 }     UP      8t1-card
```

The system labeled TNT-02 contains the following slot cards:

```
admin> show
Shelf 1 ( standalone ):
  { shelf-1 slot-1 0 }      UP      t3-card
  { shelf-1 slot-2 0 }      UP      t3-card
  { shelf-1 slot-3 0 }      UP      4ether2-card
  { shelf-1 slot-4 0 }      UP      4ether2-card
  { shelf-1 slot-5 0 }      UP      hdlc2-card
  { shelf-1 slot-6 0 }      UP      hdlc2-card
  { shelf-1 slot-7 0 }      UP      hdlc2-card
  { shelf-1 slot-8 0 }      UP      hdlc2-card
```

For information about configuring the WAN lines in each system, see “Telco considerations related to stacking” on page 43 and the *MAX TNT Hardware Installation Guide*.

For information about configuring LAN interfaces, see the *MAX TNT Network Configuration Guide*.

Sample stacking control interface configurations

In this example, the stacking control interface is on the 10.1.1.0 subnet in OSPF area 1, and the system address is set to the address of the stacking control interface. Following is the relevant IP-Interface profile on TNT-01:

```
admin> get ip-interface { { 1 3 1 } 0 }
[in IP-INTERFACE/{ { shelf-1 slot-3 1 } 0 }]
interface-address* = { { shelf-1 slot-3 1 } 0 }
ip-address = 10.1.1.1/24
proxy-mode = Off
rip-mode = routing-off
route-filter = ""
rip2-use-multicast = yes
ospf = { yes 0.0.0.1 normal 10 40 5 simple ascend0 0 1 16777215 typ+
multicast-allowed = no
multicast-rate-limit = 100
multicast-group-leave-delay = 0
directed-broadcast-allowed = yes
vrouter = ""
management-only-interface = no
```

The following commands set the TNT-01 system address to the stacking control interface:

```
admin> read ip-global
IP-GLOBAL read

admin> set system-ip-addr = 10.1.1.1

admin> write
IP-GLOBAL written
```

Following is the relevant IP-Interface profile on TNT-02:

```
admin> get ip-interface { { 1 4 1 } 0 }
[in IP-INTERFACE/{ { shelf-1 slot-4 1 } 0 }]
interface-address* = { { shelf-1 slot-4 1 } 0 }
ip-address = 10.1.1.2/24
proxy-mode = Off
rip-mode = routing-off
route-filter = ""
rip2-use-multicast = yes
ospf = { yes 0.0.0.1 normal 10 40 5 simple ascend0 0 1 16777215 typ+
multicast-allowed = no
multicast-rate-limit = 100
multicast-group-leave-delay = 0
directed-broadcast-allowed = yes
vrouter = ""
management-only-interface = no
```

The following commands set the TNT-02 system address to the stacking control interface:

```
admin> read ip-global
IP-GLOBAL read

admin> set system-ip-addr = 10.1.1.2
```

```
admin> write
IP-GLOBAL written
```

Sample stacking data interface configurations

In this example, the stacking data interface is on the 10.20.20.0 subnet in OSPF area 1. Following is the relevant IP-Interface profile on TNT-01:

```
admin> get ip-interface { { 1 3 4 } 0 }
[in IP-INTERFACE/{ { shelf-1 slot-3 4 } 0 }]
interface-address* = { { shelf-1 slot-3 4 } 0 }
ip-address = 10.20.20.1/24
proxy-mode = Off
rip-mode = routing-off
route-filter = ""
rip2-use-multicast = yes
ospf = { yes 0.0.0.1 normal 10 40 5 simple ascend0 0 1 16777215 typ+
multicast-allowed = no
multicast-rate-limit = 100
multicast-group-leave-delay = 0
directed-broadcast-allowed = yes
vrouter = ""
management-only-interface = no
```

Following is the relevant IP-Interface profile on TNT-02:

```
admin> get ip-interface { { 1 4 4 } 0 }
[in IP-INTERFACE/{ { shelf-1 slot-4 4 } 0 }]
interface-address* = { { shelf-1 slot-4 4 } 0 }
ip-address = 10.20.20.2/24
proxy-mode = Off
rip-mode = routing-off
route-filter = ""
rip2-use-multicast = yes
ospf = { yes 0.0.0.1 normal 10 40 5 simple ascend0 0 1 16777215 typ+
multicast-allowed = no
multicast-rate-limit = 100
multicast-group-leave-delay = 0
directed-broadcast-allowed = yes
vrouter = ""
management-only-interface = no
```

Configuring the stack

The following commands configure a Stacking profile on TNT-01:

```
admin> read stacking
STACKING read

admin> set enabled = yes

admin> set name = stack

admin> set multicast-interface-ip-address = 10.1.1.1

admin> set data-ip-address = 10.20.20.1

admin> list
[in STACKING (new)(changed)]
enabled = yes
name = stack
```

```
udp-port = 5150
multicast-address = 239.192.74.72
multicast-interface-ip-address = 10.1.1.1
data-ip-address = 10.20.20.1

admin> write
STACKING written
```

The Multicast-Interface-IP-Address setting specifies the stacking control interface, and the Data-IP-Address setting specifies the stacking data interface address. Note that all members of a stack must specify the same name, UDP port, and multicast address. The following commands configure a Stacking profile on TNT-02:

```
admin> read stacking
STACKING read

admin> set enabled = yes

admin> set name = stack

admin> set multicast-interface-ip-address = 10.1.1.2

admin> set data-ip-address = 10.20.20.2


admin> list
[in STACKING (new)(changed)]
enabled = yes
name = stack
udp-port = 5150
multicast-address = 239.192.74.72
multicast-interface-ip-address = 10.1.1.2
data-ip-address = 10.20.20.2

admin> write
STACKING written
```

Example of stacking a MAX TNT and MAX unit

The following example shows how to stack a MAX TNT with a MAX unit. In this example, the MAX TNT uses a single 100BaseT Ethernet as the interface for both stacking control packets and stacking data packets, which allows it to operate in a stack with the MAX unit without the use of an external router. The stack is shown in Figure 8:

Figure 8. Example stack of a MAX TNT and MAX unit



For information about configuring the WAN lines in each peer, see “Telco considerations related to stacking” on page 43 and the relevant *Hardware Installation Guide*.

For information about configuring LAN interfaces, see the *MAX TNT Network Configuration Guide*. For details about configuring stacking in the MAX, see the *MAX Network Configuration Guide*.

Sample MAX TNT IP interface configurations

In this example, the stacking control interface and the stacking data interface are both on the 10.20.20.0 subnet. OSPF is enabled and configured in area 1. In addition, the system address is set to the address of the stacking interface. Following is a sample IP-Interface profile on the MAX TNT:

```
admin> get ip-interface { { 1 3 4 } 0 }
[in IP-INTERFACE/{ { shelf-1 slot-3 4 } 0 }]
interface-address* = { { shelf-1 slot-3 4 } 0 }
ip-address = 10.20.20.1/24
proxy-mode = Off
rip-mode = routing-off
route-filter = ""
rip2-use-multicast = yes
ospf = { yes 0.0.0.1 normal 10 40 5 simple ascend0 0 1 16777215 typ+
multicast-allowed = no
multicast-rate-limit = 100
multicast-group-leave-delay = 0
directed-broadcast-allowed = yes
vrouter = ""
management-only-interface = no
```

The following commands set the system address to the stacking interface:

```
admin> read ip-global
IP-GLOBAL read

admin> set system-ip-addr = 10.20.20.1

admin> write
IP-GLOBAL written
```

Sample MAX IP interface configuration

In this example, the MAX unit is also on the 10.20.20.0 subnet. OSPF is enabled and configured in area 1.

- 1 Open Ethernet > Mod Config > Ether Options and set the MAX unit's IP address. For example:

```
90-A** Mod Config
Ether Options...
>IP Adrs=10.20.20.2/24
2nd Adrs=0.0.0.0/0
RIP=Off
RIP2 Use Multicast=No
Ignore Def Rt=Yes
Proxy Mode=Off
Filter=0
IPX Frame=N/A
```

- 2 Open Ethernet > Mod Config > OSPF Options and enable OSPF.

```
90-A** Mod Config
OSPF options...
>RunOSPF=Yes
Area=0.0.0.0
AreaType=Normal
HelloInterval=10
DeadInterval=40
```

```
Priority=5
AuthType=Simple
AuthKey=ascend0
Cost=1
ASE-type=N/A
ASE-tag=N/A
TransitDelay=1
RetransmitInterval=5
```

3 Configure the OSPF area number. For example:

```
90-A** Mod Config
OSPF options...
  RunOSPF=Yes
>Area=0.0.0.1
  AreaType=Normal
  HelloInterval=10
  DeadInterval=40
  Priority=5
  AuthType=Simple
  AuthKey=ascend0
  Cost=1
  ASE-type=N/A
  ASE-tag=N/A
  TransitDelay=1
  RetransmitInterval=5
```

4 Write the Ethernet profile.

Configuring the stack

Following is an example of configuring the Stacking profile in the MAX TNT:

```
admin> read stacking
STACKING read

admin> set enabled = yes

admin> set name = stack

admin> set multicast-address = 224.0.1.1

admin> set multicast-interface-ip-address = 10.20.20.1

admin> set data-ip-address = 10.20.20.1

admin> list
[in STACKING (new)(changed)]
enabled = yes
name = stack
udp-port = 5150
multicast-address = 224.0.1.1
multicast-interface-ip-address = 10.20.20.1
data-ip-address = 10.20.20.1

admin> write
STACKING written
```

Following is an example of configuring the Stack Options subprofile in the MAX unit. Note that all members of a stack must specify the same name, UDP port, and multicast address.

1 Open the Ethernet > Mod Config > Stack Options subprofile and enable stacking.

```
90-A** Mod Config
Stack Options...
```

```
>Stacking Enabled=Yes
Stack Name=""
UDP Port=5150
Multicast Addr=239.192.74.7
```

- 2 Set the Stack Name parameter to the same name configured in the MAX TNT.

```
90-A** Mod Config
Stack Options...
Stacking Enabled=Yes
>Stack Name=stack
UDP Port=5150
Multicast Addr=239.192.74.7
```

- 3 Leave the default UDP port number (as in the MAX TNT configuration).

- 4 Specify the same multicast address used by the MAX TNT.

```
90-A** Mod Config
Stack Options...
Stacking Enabled=Yes
Stack Name=stack
UDP Port=5150
>Multicast Addr=224.0.1.1
```

- 5 Write the Ethernet profile.

Location:

Authentication, Authorization, and Accounting (AAA)

Sharing profiles on a per-user basis

With MAX TNT TAOS 8.0.0, you can enable shared profiles on a per-connection basis even though they have been disallowed system-wide. Previously, this functionality was available only in RADIUS profiles via the Ascend-Shared-Profile-Enable attribute. Following is the new parameter, shown with its default setting:

```
[in CONNECTION/" "]
shared-prof = no
```

Parameter	Specifies
Shared-Prof	Enable/disable multiple callers to share the Connection profile, provided that IP address conflicts do not result. With the default setting of no, the setting of the Shared-Prof parameter in the IP-Global profile allows or disallows shared profiles system-wide.

If the IP-Global profile sets Shared-Prof to yes, the Shared-Prof setting in a Connection profile has no effect. However, if the IP-Global profile sets Shared-Prof to no, and a Connection profile sets it to yes, the setting in a Connection profile takes precedence. For example, with the following settings, multiple callers can call in and authenticate the Connection profile named shared-1:

```
admin> get ip-global shared-prof
[in IP-GLOBAL:shared-prof]
shared-prof = no
```

```
admin> read connection shared-1
CONNECTION/shared-1 read

admin> set shared-prof = yes

admin> set ip-options ip-routing-enabled = no

admin> write
CONNECTION/shared-1 written
```

New settings for CLID-Auth-Mode

The CLID-Auth-Mode parameter supports new `clid-first` and `dnis-first` settings in addition to the `clid-prefer` and `dnis-prefer` settings.

If CLID-Auth-Mode is set to `clid-first` or `dnis-first` and the calling-line ID (CLID) or called number (DNIS) is sent by the telco switch, the MAX TNT uses it to authenticate the call. If that level of authentication fails for any reason, or if the telco switch does not provide the calling-line ID or called number, the MAX TNT does not drop the call, but allows negotiations to proceed to password authentication. The following commands set CLID-Auth-Mode to DNIS-First:

```
admin> read answer
ANSWER-DEFAULTS read

admin> set clid-auth-mode = dnis-first

admin> write
ANSWER-DEFAULTS written
```

Extensions to DNIS management

As in earlier releases, you can use NavisAccess™ to set a Dialed Number Information Service (DNIS) number to a disabled state, which causes the MAX TNT to refuse calls to that number. This functionality is used to implement classes of service. For example, an ISP's modem pool might specify different DNIS numbers for different classes of service. When the modems reach a preset low limit, NavisAccess sets a DNIS for the low class of service into the disabled state, causing the system to refuse calls on that DNIS, while continuing to accept calls on a DNIS related to higher classes of service.

In this release, the following changes related to disabled DNIS have been made to the service MIB (`srvcgmt.mib`):

- The `dnisMgmtGlobalTable` has a new object for defining a timeout for disabled DNIS.
- Each entry in the `dnisMgmtGlobalTable` is now indexed by DNIS number, rather than by an integer.

In addition, an accounting Stop message is generated when a call is rejected because of a disabled DNIS.

Timeout for a disabled DNIS

A new timeout object is defined for limiting the disabled state. If the available resource is still below its low limit when the timeout expires, NavisAccess can set the timeout again. (NavisAccess can set the timeout at any time, even if the current timeout has not expired.) If

the resource is not below its low limit when the timeout expires, the DNIS returns to the enabled state and starts accepting incoming calls. Following is the new object:

```
dnisGlobalStatusTimeout OBJECT-TYPE
    SYNTAX  TimeTicks
    ACCESS  read-write
    STATUS  mandatory
    DESCRIPTION
        "This timeout is applied to the dnis if its dnisGlobalStatus
        is set to disabled(2). After the timeout expires, this
        dnis changes its dnisGlobalStatus to enabled(1).
        This timeout may be set at any time, even if the previous
        one has not yet expired. If the timeout is not specified by
        user at the time dnisGlobalStatus set to disabled(2), it is
        set to its default value.
        Note: When being read, this attribute shows the time interval
        remaining till expiration."
    DEFVAL { 360000 }
    ::= { dnisMgmtGlobalEntry 7 }
```

This object specifies a time interval in ticks (0.01 seconds) for a DNIS number to stay in the disabled state. After the timeout expires, the DNIS starts to accept incoming calls again, unless NavisAccess resets the timeout. The default value for the timeout is 1 hour.

New index for DNIS management table entries

In previous releases, each entry in the dnisMgmtGlobalTable was identified by an integer index. Operators had to retrieve the entire table to determine the index associated with a particular DNIS before modifying a DNIS status. This resulted in multiple Simple Network Management Protocol (SNMP) protocol data units (PDUs), increasing traffic and response time.

In this release, dnisMgmtGlobalTable entries are identified by dnisGlobalPhoneNumber, which is a DisplayString (a character string) of from 4 to 24 characters. The index to a particular entry in the table contains from 5 to 25 subindexes, because the first subindex contains the number of subindexes (the number of digits in the telephone number). For example, if the DNIS number is 1234, the object ID for the dnisGlobalStatus is as follows:

```
dnisGlobalStatus.4.49.50.51.52.
```

The value 4 is the number of digits in the DNIS, and 49, 50, 51, and 52 are decimal representations of the ASCII values for 1, 2, 3, and 4 (0x31, 0x32, 0x33, and 0x34).

The part of the DNIS management table related to the modified index follows:

```
INDEX { dnisGlobalPhoneNumber }
::= { dnisMgmtGlobalTable 1 }

DnisMgmtGlobalEntry ::=
    SEQUENCE {
        dnisGlobalIndex          INTEGER,
        dnisGlobalPhoneNumber    DisplayString,
        dnisGlobalStatus         INTEGER,
        dnisGlobalCallsAccepted  Counter,
        dnisGlobalCallsDropped  Counter,
        dnisGlobalAction         INTEGER,
        dnisGlobalStatusTimeout  TimeTicks
    }
```



```
dnisGlobalIndex OBJECT-TYPE
    SYNTAX INTEGER
    ACCESS read-only
    STATUS deprecated
    DESCRIPTION
        "A unique value, greater than zero, for each
        DNIS telephone number. The value ranges from 1 to
        'dnisMgmtGlobalNumEntries' and identifies which
        DNIS telephone number is associated with.
        It is recommended that the value are assigned
        continuously starting from 1. The value must remain
        constant at least from one re-initialization
        of the entity's network management system to the next
        re-initialization.
        Note: This attribute is no longer used as an index
        in this table. It always returns 1."
 ::= { dnisMgmtGlobalEntry 1 }

dnisGlobalPhoneNumber OBJECT-TYPE
    SYNTAX DisplayString (SIZE(4..24))
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION
        "4 to 24 digits number. The DNIS is the telephone number
        dialed to access the hunt group.
        This number is extracted from incoming calls or from
        dnisGlobalActioni. Note that if this number contains
        characters other than digits, the agent returns error. "
 ::= { dnisMgmtGlobalEntry 2 }
```

Accounting Stop message due to a disabled DNIS

An accounting Stop message is now generated when an incoming call is rejected because of a disabled DNIS number. In the Stop message, the value of the Ascend-Disconnect-Cause is 370 and the value of the Ascend-Connect-Progress attribute is 11. These are new values introduced for calls rejected on the basis of DNIS management.

Disconnect Code	Meaning
370	DNIS Denied.

Progress Code	Meaning
11	Dialed Service Blocked.

For example:

```
NAS-Port = 0
...
Acct-Status-Type = Stop
Acct-Delay-Time = 0
...
Ascend-Disconnect-Cause = 370
Ascend-Connect-Progress = 11
Ascend-Xmit-Rate = 0
Ascend-Data-Rate = 0
Ascend-PreSession-Time = 0
```

```
Ascend-Pre-Input-Octets = 0
Ascend-Pre-Output-Octets = 0
Ascend-Pre-Input-Packets = 0
Ascend-Pre-Output-Packets = 0
```

Bidirectional CHAP

In previous releases, a MAX TNT unit performed only unidirectional authentication between PPP devices. The called device authenticated the calling device. The calling device could not determine whether the remote site was the correct one. With MAX TNT TAOS 8.0.0, you can set up bidirectional Challenge Handshake Authentication Protocol (CHAP) authentication between the calling PPP device and the called PPP device.

Bidirectional CHAP increases compliance with the RFC 1994 standard for PPP CHAP authentication. Note that the feature is not implemented for PAP-based authentication (PAP, PAP-TOKEN, or PAP-TOKEN-CHAP).

Note: As noted in RFC 1994, a security hole can occur when you use bidirectional authentication for an incoming call if the secrets used in both directions are identical. Bidirectional authentication in TAOS has been developed to avoid the security hole, even if the secrets are identical. For best results, however, Lucent recommends that you specify a different secret for each authentication direction.

Bidirectional CHAP is supported locally and through RADIUS.

Configuring bidirectional CHAP on a MAX TNT unit


The following sections describe how to configure bidirectional CHAP in local profiles. You can choose one or more of the following configurations:

- Setting up bidirectional CHAP for all incoming calls
- Setting up bidirectional CHAP for selected incoming calls
- Setting up bidirectional CHAP for outgoing calls

Setting up bidirectional CHAP for all incoming calls

Figure 9 shows a configuration in which a MAX TNT and its dial-in clients authenticate each other by means of bidirectional CHAP. One or more clients can dial into the MAX TNT unit. The MAX TNT unit authenticates each calling device by means of a Connection profile, and each dial-in client authenticates the MAX TNT by means of the Send-Password value.

Figure 9. Bidirectional CHAP for all incoming calls to the MAX TNT unit



To configure bidirectional CHAP on the MAX TNT unit for all incoming calls, proceed as follows:

- 1 Make Answer-Defaults the working profile.
- 2 List the PPP-Answer subprofile.
- 3 Set Receive-Auth-Mode to `any-ppp-auth` or `chap-ppp-auth`.
- 4 Set Bi-Directional-Auth to `required` or `allowed`. Required specifies that bidirectional authentication must be carried out or the call is dropped. Allowed specifies that authentication *can* be bidirectional. The MAX TNT unit identifies the calling device, and the calling device can identify the MAX TNT unit, but the calling device need not do so for the call to be accepted.
- 5 Write the Answer-Defaults profile.
- 6 For each incoming call, create or read a Connection profile, and make it the working profile.
- 7 List the PPP-Options subprofile.
- 8 Set Send-Password to any text string. The password you specify is the one sent to the calling unit during the authentication initiated by the calling unit.
- 9 Set Recv-Password to any text string. The password you specify is the one sent by the calling unit during the authentication initiated by the MAX TNT unit.
- 10 Write the Connection profile.

Note: When the Receive-Auth-Mode parameter is set to `any-ppp-auth`, the MAX TNT unit can accept both Password Authentication Protocol (PAP) and CHAP authentication. The Bi-Directional-Auth setting is used only if a form of CHAP authentication has been negotiated during Link Control Protocol (LCP) negotiation. If any form of PAP authentication has been negotiated, and Bi-Directional-Auth is set to `required`, the MAX TNT unit authenticates the calling unit, and authentication takes place in one direction only.

The following is a sample configuration:

```
admin> read answer-defaults
ANSWER-DEFAULTS read

admin> set ppp-answer receive-auth-mode = chap-ppp-auth
admin> set ppp-answer bidirectional-auth = required
admin> write
ANSWER-DEFAULTS written


admin> read connection robin
CONNECTION/robin read

admin> set ppp-options send-password = sendpw
admin> set ppp-options recv-password = recvpw
admin> write
CONNECTION/robin written
```

Setting up bidirectional CHAP for selected incoming calls

Figure 10 shows a configuration in which the MAX TNT unit authenticates the calling device by means of Calling-Line ID (CLID) or Dialed Number Information Service (DNIS) authentication. The dial-in client and the MAX TNT then authenticate each other by means of CHAP.

Figure 10. Bidirectional CHAP for selected calls



To configure bidirectional CHAP on the MAX TNT unit for selected incoming calls, proceed as follows:

- 1 Make Answer-Defaults the working profile.
- 2 Set Profiles-Required to *yes*.
- 3 Set CLID-Auth-Mode to *clid-require*, *clid-prefer*, *dnis-require*, or *dnis-prefer*.
- 4 List the PPP-Answer subprofile.
- 5 Set Bi-Directional-Auth to *none* or *allowed*.
- 6 Write the Answer-Defaults profile.
- 7 Select or create the Connection profile for which you want to set up bidirectional CHAP, and make it the working profile.
- 8 If CLID-Auth-Mode is set to *clid-require* or *clid-prefer*, set the CLID value to the CLID.
- 9 If CLID-Auth-Mode is set to *dnis-require* or *dnis-prefer*, set the CalledNumber value to the number the calling party dials.
- 10 List the PPP-Options subprofile.
- 11 Set Send-Password to any text string. The password you specify is the one sent to the calling unit during the authentication initiated by the calling unit.
- 12 Set Recv-Password to any text string. The password you specify is the one sent by the calling unit during the authentication initiated by the MAX TNT unit.
- 13 Set Send-Auth-Mode to *chap-ppp-auth*. This value indicates the mode for both incoming and outgoing authentication.
- 14 Set Bi-Directional-Auth to *required* or *allowed*. Required specifies that bidirectional authentication must be carried out or the call is dropped. Allowed specifies that authentication *can* be bidirectional. The MAX TNT unit identifies the calling device, and the calling device can identify the MAX TNT unit, but the calling device need not do so for the call to be accepted.
- 15 Write the Connection profile.

The following is a sample configuration:

```
admin> read answer-defaults
ANSWER-DEFAULTS read

admin> set profiles-required = yes

admin> set clid-auth-mode = clid-require

admin> set ppp-answer bidirectional-auth = allowed
```

```
admin> write
ANSWER-DEFAULTS written

admin> read connection robin
CONNECTION/robin read

admin> set clid = 1234567

admin> set ppp-options send-password = "passin"
admin> set ppp-options recv-password = "passout"
admin> set ppp-options send-auth-mode = chap-ppp-auth
admin> set ppp-options bi-directional-auth = allowed

admin> write
CONNECTION/robin written
```

Setting up bidirectional CHAP for outgoing calls

To set up bidirectional CHAP on the MAX TNT unit for outgoing calls, proceed as follows:

- 1 Make the Connection profile the working profile.
- 2 List the PPP-Options subprofile.
- 3 Set Send-Auth-Mode to chap-ppp-auth, cache-token-ppp-auth, or ms-chap-ppp-auth. If you specify any other authentication mode, bidirectional authentication does not take place, even if Bi-Directional-Auth is set to allowed or required.
- 4 Set Bi-Directional-Auth to required or allowed. Required specifies that bidirectional authentication must be carried out or the call is dropped. Allowed specifies that authentication *can* be bidirectional. The MAX TNT unit identifies the called device if the called device accepts the authentication, and the called device can identify the MAX TNT unit, but the called device need not do so for the call to be accepted.
- 5 Set Send-Password to a text string specifying the password sent to the called device during the authentication initiated by the MAX TNT unit.
- 6 Set Recv-Password to a text string specifying the password sent by the called unit during the authentication initiated by the called unit.
- 7 Set Substitute-Recv-Name to a text string. The called party's name is compared against the value you specify. If the called party's name is different, the call is not established. If you do not specify a value for Substitute-Recv-Name, the called party's name is compared against the dialout profile name.
- 8 Write the Connection profile.

The following is a sample configuration:

```
admin> read connection robin
CONNECTION/robin read

admin> set ppp-options send-auth-mode = chap-ppp-auth
admin> set ppp-options bi-directional-auth = required
admin> set ppp-options send-password = sendpw
admin> set ppp-options recv-password = recvpw
admin> set ppp-options substitute-recv-name = subname

admin> write
CONNECTION/robin written
```

Parameter reference entries

Bi-Directional-Auth

Description: Specifies whether CHAP authentication must be bidirectional.

Usage: Specify none (the default), allowed, or required.

- None (the default) specifies that authentication is unidirectional. The called device identifies the calling one. The MAX TNT unit prevents the authentication in which the calling party identifies the called party.

- Allowed specifies that authentication can be bidirectional.

When the MAX TNT unit is the called device, it identifies the calling device. The system also allows the calling device to authenticate the MAX TNT unit, but this authentication is not mandatory. Therefore, if the calling device does not authenticate the MAX TNT unit, the MAX TNT unit can still accept the call.

When the MAX TNT unit is the calling device, it answers the authentication initiated by the called device. The MAX TNT unit tries to negotiate authentication in the opposite direction as well, but if the called device refuses this second authentication option, the call is still established.

- Required specifies that authentication must be bidirectional. The MAX TNT unit requires that both the calling and called devices authenticate each other. If authentication is not performed in both directions, the MAX TNT unit rejects the call (in the case of an incoming call) or tears down the call (in the case of an outgoing call).

Example: `set bi-directional-auth = allowed`

Dependencies: Consider the following:

- If you specify `allowed` or `required`, and the second authentication is attempted, it must be successful. Otherwise, the MAX TNT unit rejects the call (in the case of an incoming call) or tears down the call (in the case of an outgoing call).
- Bidirectional authentication is applicable only if the authentication mode is CHAP, MS-CHAP or CACHE-TOKEN.
- When Receive-Auth-Mode is set to `any` or `either`, and PAP authentication is negotiated, bidirectional authentication is automatically disabled, even if the Bi-Directional-Auth is set to `required`. For example, suppose you set Receive-Auth-Mode to Any-PPP-Auth and Bi-Directional-Auth to Required. If an incoming call occurs and the authentication negotiated is PAP, the authentication takes place in one direction only.

Location: Answer-Defaults > PPP-Answer, Connection > PPP-Options

See Also: Substitute-Recv-Name

Substitute-Recv-Name

Description: Specifies the PPP called device's name during outgoing calls. Because bidirectional authentication provides a way to formally authenticate the called device during an outgoing call, the name of the device must be checked against a locally defined name. The name can be the dialout profile name or a substituted name.

Usage: Specify a string of up to 23 characters. The default is null.

Example: `set substitute-recv-name = fred`

Dependencies: Consider the following:

- The value you specify for Substitute-Recv-Name is used only during outgoing calls that use bidirectional authentication.
- If you accept the default of null for Substitute-Recv-Name, the name of the called device is checked against the dialout profile name.
- Substitute-Recv-Name allows an additional RADIUS lookup during an outgoing call.
- Because Substitute-Recv-Name represents the called device's real name, it is sent in RADIUS accounting Start and Stop messages.

Location: Connection > PPP-Options

See Also: Bi-Directional-Auth

Configuring bidirectional CHAP in RADIUS

The following sections describe how to configure bidirectional CHAP in RADIUS. You can use one of the following configurations:

- Setting up bidirectional CHAP in RADIUS for incoming calls
- Setting up bidirectional CHAP in RADIUS for outgoing calls
- Setting up selective bidirectional CHAP with callback
- Setting up an outgoing call with double RADIUS lookups

Setting up bidirectional CHAP in RADIUS for incoming calls

You can configure selective bidirectional authentication by using CLID or DNIS preauthentication in a pseudo-user profile and then specifying two passwords in the user profile.

In the pseudo-user profile, specify CLID or DNIS authentication and then set the Ascend-Bi-Directional-Auth attribute to Bi-Directional-Auth-Allowed or Bi-Directional-Auth-Required:

- Bi-Directional-Auth-Allowed specifies that authentication can be bidirectional. The MAX TNT unit identifies the calling device. The system also allows the calling device to authenticate the MAX TNT unit, but this authentication is not mandatory. Therefore, if the calling device does not authenticate the MAX TNT unit, the MAX TNT unit can still accept the call.
- Bi-Directional-Auth-Required specifies that authentication must be bidirectional.

In the following pseudo-user profile, bidirectional authentication is required:

```
111886067 Password = "Ascend-CLID"
    Service-Type = Framed,
    Ascend-Require-Auth = Require-Auth,
    Ascend-Auth-Type = Auth-CHAP,
    Ascend-Send-Auth = Send-Auth-CHAP,
    Ascend-Bi-Directional-Auth = Bi-Directional-Auth-Required
```

In the user profile, Ascend-Send-Secret is set to the password sent to the called device during the authentication initiated by the MAX TNT unit:

```
Mikel Password = "passin"
    Service-Type = Framed,
    Ascend-Send-Secret = "passout",
```

```
Framed-Protocol = PPP
Framed-IP-Address = 111.5.1.1
Framed-IP-Netmask = 255.255.255.255
Ascend-Data-Svc = Switched-64K
Ascend-Route-IP = Route-IP-Yes
```

Note that the Answer-Defaults profile must contain the desired bidirectional authentication mode (none, required, or allowed) if CLID or DNIS preauthentication is not in use. The pseudo-user profile can be suppressed (unused), and the user profile must contain the Ascend-Bi-Directional-Auth attribute.

Setting up bidirectional CHAP in RADIUS for outgoing calls

To configure a RADIUS dialout profile that makes use of bidirectional authentication, proceed as follows:

- 1 Set User-Name to the name of the called party, and Password to ascend.
- 2 Set Ascend-Send-Auth to send-auth-chap.
- 3 Set Ascend-Send-Secret to the text of the secret sent to the called device.
- 4 Set Ascend-Receive Secret to the text of the secret received from the called device.
- 5 Set Ascend-Bi-Directional-Auth to bi-directional-auth-allowed or bi-directional-auth-required.
- 6 Set Ascend-Recv-Name to the name of the called party.

For example:

```
Mikel-out Password = "ascend"
  Service-Type = Outbound,
  User-Name = "Mikel",
  Framed-Protocol = PPP,
  Framed-IP-Address = 111.5.1.1,
  Framed-IP-Netmask = 255.255.255.0,
  Ascend-Dial-Number = 90492386067,
  Ascend-Data-Svc = Switched-64K,
  Ascend-Send-Auth = Send-Auth-CHAP,
  Ascend-Send-Secret = "passout",
  Ascend-Receive-Secret = "passin",
  Ascend-Bi-Directional-Auth = Bi-Directional-Auth-Required,
  Ascend-Route-IP = 1

route-tnt-pat-1 Password = "ascend"
  Service-Type = Outbound,
  Framed-Route = "111.5.1.0/30 111.5.1.1 1 n Mikel-out"
```

Setting up selective bidirectional CHAP with callback

To configure bidirectional CHAP with callback, you must carry out the following steps:

- Create a first-tier pseudo-user profile.
- Create a second-tier user profile.

In the first-tier pseudo-user profile, proceed as follows:

- 1 Set User-Name to the name of the called party, and Password to ascend.
- 2 Set Ascend-Require-Auth to require-auth.

- 3 Set Ascend-Send-Auth to send-auth-chap.
- 4 Set Ascend-Bi-Directional-Auth to bi-directional-auth-allowed or bi-directional-auth-required.

For a global bidirectional CHAP callback, the first-tier pseudo-user profile is not used. In the second-tier user profile, proceed as follows:

- 1 Set Ascend-Send-Auth to send-auth-chap.
- 2 Set Ascend-Bi-Directional-Auth to bi-directional-auth-allowed or bi-directional-auth-required.
- 3 Set Ascend-Callback to callback=yes.

The following example shows the configuration required for callback. In the first-tier pseudo-user profile, bidirectional authentication is selectively determined during DNIS preauthentication, and the system performs bidirectional authentication for both incoming and outgoing calls. The second-tier user profile is configured for bidirectional CHAP with callback.

```
8940 Password = "Ascend-DNIS"
    Service-Type = Outbound,
    Ascend-Require-Auth = Require-Auth,
    Ascend-Auth-Type = Auth-CHAP,
    Ascend-Send-Auth = Send-Auth-CHAP,
    Ascend-Bi-Directional-Auth = Bi-Directional-Auth-Required

Mikel_cb Password = "passin"
    Service-Type = Framed,
    Ascend-Send-Secret = "pass",
    Framed-Protocol = MP,
    Ascend-Base-Channel-Count = 2,
    Ascend-Minimum-Channels = 1,
    Ascend-Maximum-Channels = 2,
    Framed-IP-Address = 111.5.1.1,
    Framed-IP-Netmask = 255.255.255.255,
    Ascend-Dial-Number = 90492386067,
    Ascend-Data-Svc = Switched-64K,
    Ascend-Send-Auth = Send-Auth-CHAP,
    Ascend-Bi-Directional-Auth = Bi-Directional-Auth-Required,
    Ascend-Callback = Callback-Yes,
    Ascend-Callback-Delay = 10,
    Ascend-Route-IP = 1
```

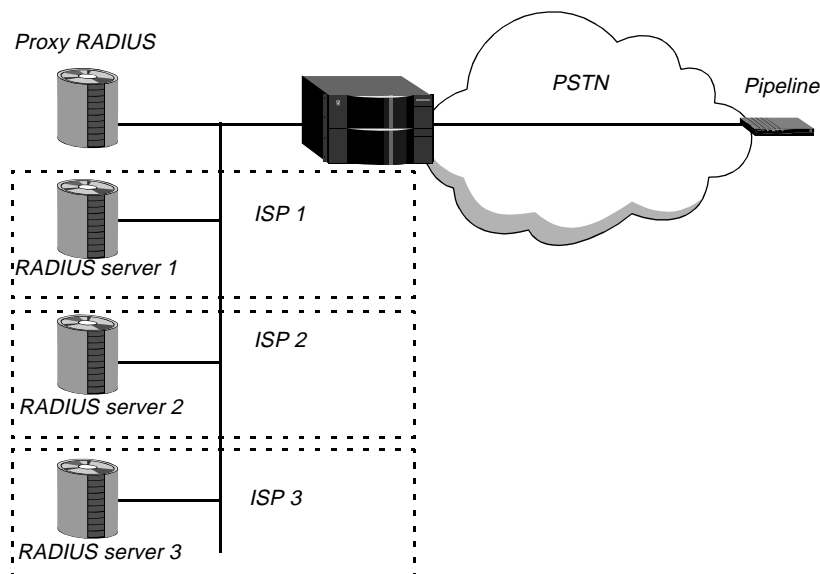
Setting up an outgoing call with double RADIUS lookups

This section discusses the following:

- The circumstances under which you might use double RADIUS lookups
- The procedure for setting up RADIUS lookups
- The message sequence during RADIUS lookups

In larger networks, several ISPs might be hosted on a single physical network, such as the one shown in Figure 11. Each ISP typically has its own RADIUS server, while the network provider uses a proxy RADIUS server. The MAX TNT unit interacts only with proxy RADIUS server. The proxy server can answer some requests locally, and forward other requests to the RADIUS server of an ISP. Typically, an ISP requires that all of its users be authenticated by its own RADIUS server, and not by the network provider's equipment.

Figure 11. Bidirectional CHAP in a multiprovider network



During an outgoing call with bidirectional authentication, the MAX TNT unit first recovers the dialout profile. Once the call is brought up, the MAX TNT unit must authenticate the called party, in this case a Pipeline unit. The authentication decision must be made by the ISP's RADIUS server, requiring a second RADIUS lookup.

When you set up double RADIUS lookups, the dialout profile is split into two profiles—the first-tier dialout profile and the second-tier user profile. The dialout profile contains all dialout parameters needed to establish the outgoing call, and the user profile contains information for authenticating the called device.

Consider the following first-tier dialout profile, configured for bidirectional CHAP authentication:

```
pipe-pat-out Password = "ascend"
  Service-Type = Outbound,
  Framed-Protocol = PPP,
  Framed-IP-Address = 10.4.8.8,
  Framed-IP-Netmask = 255.255.255.0,
  Ascend-Dial-Number = 90492386067,
  Ascend-Data-Svc = Switched-64K,
  Ascend-Send-Auth = Send-Auth-CHAP,
  Ascend-Send-Secret = "passin",
  Ascend-Bi-Directional-Auth = Bi-Directional-Auth-Required,
  Ascend-Recv-Name = "pipe-pat",
  Ascend-Route-IP = 1
```

To enforce the second RADIUS lookup, the dialout profile name (pipe-pat-out in this example) must be different from the name of the called device in the user profile. The Ascend-Recv-Name attribute specifies the name of the called device, in this case pipe-pat.

In the following second-tier user profile, the called party's name is pipe-pat and the receive-password is pass.

```
pipe-pat Password = "pass"
  Service-Type = Framed
  Ascend-Route-IP = 1"
```

You can disable the double RADIUS lookup by naming the dialout profile with the peer's name and by omitting the Ascend-Recv-Name attribute. Use the User-Name attribute to rename the profile (in this case to pipe-pat):

```
pipe-pat-out Password = "ascend"
    Service-Type = Outbound,
    User-Name = "pipe-pat",
    Framed-Protocol = PPP,
    Framed-IP-Address = 10.4.8.8,
    Framed-IP-Netmask = 255.255.255.0,
    Ascend-Dial-Number = 90492386067,
    Ascend-Data-Svc = Switched-64K,
    Ascend-Send-Auth = Send-Auth-CHAP,
    Ascend-Send-Secret = "passin",
    Ascend-Bi-Directional-Auth = Bi-Directional-Auth-Required,
    Ascend-Receive-Secret = "pass",
    Ascend-Route-IP = 1
```

A call using two RADIUS lookups passes through the follow messaging sequence:

- 1 The MAX TNT unit requests a dialout profile from RADIUS.
- 2 RADIUS sends the dialout profile to the MAX TNT unit.
- 3 The MAX TNT unit makes an ISDN call to the remote device.
- 4 The ISDN call is connected.
- 5 The MAX TNT unit and the called party perform LCP exchanges.
- 6 The called party sends a challenge request to the MAX TNT unit.
- 7 The MAX TNT unit responds with a challenge response.
- 8 The called party informs the MAX TNT unit about whether the first level of authentication has been successful.
- 9 If the first authentication was successful, the MAX TNT unit sends a challenge request to the called party.
- 10 The called party responds with a challenge response.
- 11 The MAX TNT unit sends the authentication request to RADIUS, which performs the second lookup.
- 12 The RADIUS server informs the MAX TNT unit about whether the authentication was successful.
- 13 If the authentication was successful, the MAX TNT unit informs the called party that it has been authenticated.

RADIUS attribute reference entries

Ascend-Bi-Directional-Auth (46)

Description: Specifies whether CHAP authentication must be bidirectional.

Usage: Ascend-Bi-Directional-Auth appears in an Access-Accept packet. Specify Bi-Directional-Auth-None (0), Bi-Directional-Auth-Allowed (1), or Bi-Directional-Auth-Required (2).

- Bi-Directional-Auth-None (0) specifies that authentication is unidirectional. The called device identifies the calling device. The MAX TNT unit prevents the authentication in which the calling party identifies the called party.

- Bi-Directional-Auth-Allowed (1) specifies that authentication can be bidirectional.

When the MAX TNT unit is the called device, it identifies the calling device. The system also allows the calling device to authenticate the MAX TNT unit, but this authentication is not mandatory. Therefore, if the calling device does not authenticate the MAX TNT unit, the MAX TNT unit can still accept the call.

When the MAX TNT unit is the calling device, it answers the authentication initiated by the called device. The MAX TNT unit tries to negotiate authentication in the opposite direction as well, but if the called device refuses this second authentication option, the call is still established.

- Bi-Directional-Auth-Required (2) specifies that authentication must be bidirectional. The MAX TNT unit requires that both the calling and called devices authenticate each other. If authentication is not performed in both directions, the MAX TNT unit rejects the call (in the case of an incoming call) or tears down the call (in the case of an outgoing call).

Dependencies: Bidirectional authentication is applicable only if the authentication mode is CHAP, MS-CHAP, or CACHE-TOKEN. If you specify Bi-Directional-Auth-Allowed or Bi-Directional-Auth-Required, and the second authentication is attempted, the authentication must be successful. Otherwise, the MAX TNT unit rejects the call (in the case of an incoming call) or tears down the call (in the case of an outgoing call).

See Also: Ascend-Recv-Name (45)

Ascend-Recv-Name (45)

Description: Ascend-Bi-Directional-Auth appears in an Access-Accept packet and specifies the PPP called device's name during outgoing calls. Because bidirectional authentication provides a way to formally authenticate the called device during an outgoing call, the name of the device must be checked against a locally defined name. The name can be the dialout profile name or a substituted name.

Usage: Specify a string of up to 23 characters. The default is null.

Dependencies: Consider the following:

- The value you specify for Ascend-Recv-Name is used only during outgoing calls that use bidirectional authentication.
- If you accept the default of null for Ascend-Recv-Name, the name of the called device is checked against the dialout profile name.
- Because Ascend-Recv-Name represents the called device's real name, it is sent in RADIUS accounting Start and Stop messages.

See Also: Ascend-Bi-Directional-Auth (46)

Username available for first-tier DNIS authentication

In previous releases, if only first-tier DNIS authentication was performed, no username information was available for SNMP, Syslog, or RADIUS accounting records. With MAX TNT TAOS 8.0.0, if only first-tier DNIS authentication is performed and the profile contains a User-Name attribute-value pair, the RADIUS server returns the value of the User-Name attribute in its DNIS Auth reply. If second-tier user-password authentication is performed, the username information is taken from the login name, as in previous releases.

Following is a sample DNIS-authenticated RADIUS profile that includes the User-Name attribute:

```
3735 Password = "Ascend-DNIS"
    User-Name = "johnfan",
    Service-Type = Login-User,
    Ascend-Require-Auth = Not-Require-Auth,
    Login-Service = TCP-Clear,
    Login-Host = 10.40.40.36,
    Login-TCP-Port = 7,
    Ascend-Idle-Limit = 0
```

The User-Name value is returned to the MAX TNT and is provided in the SNMP serviceChanged events and session table, Syslog messages and RADIUS Start/Stop records.

NAS-Port-Type specification for local profiles

You can now specify digital or analog service on a per-connection basis. Previously, this functionality was available only through the RADIUS NAS-Port-Type attribute (attribute 61). Following is the relevant parameter, shown with its default setting:

```
[in CONNECTION/"":telco-options]
nas-port-type = any
```

RADIUS and corresponding local profile settings

The following table shows how the local profile settings correspond to RADIUS attribute-value pairs for the RADIUS NAS-Port-Type attribute:

RADIUS settings	Corresponding local profile settings
NAS-Port-Type = Async	nas-port-type = analog or: nas-port-type = any
NAS-Port-Type = Sync	nas-port-type = digital or: nas-port-type = any
NAS-Port-Type = ISDN_Sync	nas-port-type = digital or: nas-port-type = any
NAS-Port-Type = ISDN_Async_V120	nas-port-type = digital or: nas-port-type = any
NAS-Port-Type = ISDN_Async_V110	nas-port-type = digital or: nas-port-type = any
NAS-Port-Type = Virtual	nas-port-type = any

Parameter reference entry

Nas-Port-Type

Description: Type of service for the session. The default setting enables unrestricted service. Setting the parameter to digital or analog restricts service to the specified type.

Usage: Specify one of the following values:

- any (the incoming call is routed to an analog, digital, or virtual modem)
- digital (the call is routed to a digital modem)

- analog (analog modem).

The `digital` setting restricts the profile to synchronous links, V.110 connections, and V.120 connections. The `analog` setting restricts the profile to asynchronous connections on an analog line. The `any` setting is a superset that covers the above cases as well as virtual modems.

Example: `set nas-port-type = digital`

Location: Connection > Telco-Options

RADIUS: Enhanced IETF compliance in VSA compatibility mode

In previous releases, MAX TNT units appended a null character to User-Name (1) and User-Password (2) values in Auth-Request packets. This is still the unit's default behavior when it is not operating in VSA compatibility mode. Following is the relevant parameter, shown with its default setting:

```
[in EXTERNAL-AUTH:rad-auth-client]
auth-radius-compat = old-ascend
```

To maintain compliance with the IETF RADIUS standard, MAX TNT units no longer append a null character to the User-Name (1) and User-Password (2) values in Auth-Request packets when the unit is in VSA compatibility mode. The following commands configure the unit for VSA compatibility mode:

```
admin> read external-auth
EXTERNAL-AUTH read

admin> set rad-auth-client auth-radius-compat = vendor-specific

admin> write
EXTERNAL-AUTH written
```

For details about these settings, see the *MAX TNT Reference Guide*.

RADIUS: Tunnel attribute sets with tags and preferences

The *RADIUS Attributes for Tunnel Protocol Support Internet-Draft* defines a set of RADIUS attributes designed to support transparent tunneling to dial-in networks, where a tunnel is created automatically without any explicit action by the user. To support this type of tunneling, the user's profile includes a primary attribute set, which specifies all of the values required to set up the tunnel, as well as additional attribute sets that can be used to establish a tunnel if the primary server is unavailable.

Note: Use of tunneling attribute tags and preferences requires the NavisRadius™ product or another RADIUS server that supports them.

Overview of attribute sets and tags

A *tag* is a number from 1 to 31 that you can add to one or more of the RADIUS attributes listed in "Tunnel attributes used with tags" on page 72. Attributes that share the same tag number form an attribute set. Attribute sets in the same user profile are processed in numeric order (the set with tag 1 is processed before the set with tag 2, and so forth), unless the sets are reordered by means of the Tunnel-Preference attribute.

A tag value of 0 (zero) is considered untagged. Untagged attribute sets are processed before tagged attribute sets, unless a Tunnel-Preference setting specifies otherwise.

A tag is separated from an attribute-value pair by a colon. Following is a sample profile that specifies three attribute sets, tagged 1, 2, and 3:

```
joebloggs User-Password = "murphy"
  Tunnel-Type = L2TP : 1,
  Tunnel-Server-Endpoint = "1.1.1.1" : 1,
  Tunnel-Password = "loloaagic" : 1,
  Tunnel-Type = L2TP : 3,
  Tunnel-Server-Endpoint = "3.3.3.3" : 3,
  Tunnel-Password = "i82qb4ip" : 3,
  Tunnel-Type = L2F : 2,
  Tunnel-Server-Endpoint = "2.2.2.2" : 2,
  Tunnel-Password = "itsAsecret" : 2
```

This profile specifies that the NAS (the MAX TNT) must attempt first to establish an L2TP tunnel to the LNS at 1.1.1.1. If that attempt fails, the system attempts to bring up an L2F tunnel to a server at 2.2.2.2. If that attempt also fails, the system tries an L2TP tunnel to 3.3.3.3.

In this release, a user profile can specify up to 32 tunnel attribute sets. However, because the system waits a certain interval before each attempt to initiate a tunnel and retries a certain number of times, the client's PPP connection typically times out before 32 tunnel attempts are made. For an example of tunnel timer and retry settings, see "L2TP timer options" on page 288.

Supported tunnel protocols

In this release, RADIUS attribute tags can be used for all supported tunnel protocols. The number of attribute sets used is limited for some protocols, as shown in Table 15:

Table 15. Tunnel protocols and tagged attribute sets

Tunnel protocol	Attribute sets used
L2TP	All specified attribute sets are used.
L2F	All specified attribute sets are used.
PPTP	Only the attribute set with the highest priority is used. Priority is defined by the Tunnel-Preference (83) value or by tag order.
ATMP	Only the two sets with the highest priority are used. (From the second attribute set, only the Tunnel-Server-Endpoint (67) value is used. Other values can be omitted.) Priority is defined by the Tunnel-Preference (83) value or by tag order.

In the case of L2TP and L2F, you can use the DNS list attempt feature with this feature.

All the attribute sets in a profile must specify similar tunnel protocols, either all layer 3 tunnels (such as ATMP) or layer 2 tunnels (such as L2TP or L2F). You can mix L2TP and L2F, but not with ATMP. The following examples show two valid cases:

```
JL2 User-Password = example
  Tunnel-Type = L2TP :1,
  Tunnel-Server-Endpoint = LNS-a.example.com :1,
```

```
Tunnel-Type = L2F :2,  
Tunnel-Server-Endpoint = L2FGW.example.com :2  
  
UL3 User-Password = example  
Tunnel-Type = ATMP :1,  
Tunnel-Server-Endpoint = HA-a.example.com :1,  
Tunnel-Server-Endpoint = HA-b.example.com :2,  
Tunnel-Password = HApasword :1,  
Tunnel-Private-Group-ID = MyHomeNet :1
```

Tunnel attributes used with tags

Following are the relevant tunnel attribute-value pairs:

RADIUS attribute	Value
Tunnel-Type (64)	Tunneling protocol(s) to be used. In this release, only L2TP (3) and L2F (2) currently operate with full tunnel attribute and tag support.
Tunnel-Medium-Type (65)	Medium for establishing the tunnel. Currently, IP (1) is the only supported value.
Tunnel-Server-Endpoint (67)	IP address or hostname of the tunnel endpoint. If a DNS lookup returns several IP addresses, the system attempts to establish a tunnel to each address in turn.
Tunnel-Password (69)	Shared secret for authenticating the tunnel.
Tunnel-Preference (83)	<p>Numeric preference value for an attribute set. If more than one set of tunneling attributes is returned by the RADIUS server to the MAX TNT, the Tunnel-Preference attribute can be included in a set to indicate its relative preference, with the lowest preference value designating the most preferred set.</p> <p>If no Tunnel-Preference is included in any of the attribute sets, the sets are processed in the order of their respective tag numbers.</p> <p>If some but not all attribute sets contain a Tunnel-Preference value, the attribute sets without a Tunnel-Preference are designated as the least preferred sets.</p> <p>Attribute sets with identical preferences are processed in random order.</p>
Tunnel-Client-Auth-ID (90)	Name used by the tunnel initiator (the NAS) to authenticate the tunnel server. Currently, this attribute is supported only for L2F tunnels. If it is not specified, the L2F-System-Name value in the local profile configuration is used. For details, see “Limited support for Layer 2 Forwarding (L2F)” on page 293.
Ascend-Tunnel-VRouter-Name (31)	Name of a virtual router (VRouter) to use for establishing the L2TP or L2F tunnel. The specified VRouter must exist on the LAC. For details, see “VRouter support for L2TP connections” on page 303.
Tunnel-Private-Group-ID (81)	Name of the Connection profile that defines the link on which the ATMP Home Agent transmits packets it receives from the mobile client. This attribute is supported only for ATMP tunnels. The value is used only if the Home Agent is in gateway mode. See Ascend-Home-Network-Name (185) for an alternate.

The MAX TNT currently ignores the following attributes if it receives them in a RADIUS response:

- Tunnel-Assignment-ID (82)
- Tunnel-Client-Endpoint (66)

Example of reordering sets using Tunnel-Preference

Following is a sample profile that specifies three attribute sets, tagged 1, 2, and 3, with a Tunnel-Preference value that changes the order in which the NAS attempts tunnel establishment for this user:

```
joebloggs Password = "murphy"
  Tunnel-Type = L2TP : 1,
  Tunnel-Server-Endpoint = "1.1.1.1" : 1,
  Tunnel-Password = "loloaqic" : 1,
  Tunnel-Type = L2TP : 3,
  Tunnel-Server-Endpoint = "3.3.3.3" : 3,
  Tunnel-Password = "i82qb4ip" : 3,
  Tunnel-Type = L2F : 2,
  Tunnel-Server-Endpoint = "2.2.2.2" : 2,
  Tunnel-Password = "itsAsecret" : 2,
  Tunnel-Preference = 100 : 2,
  Tunnel-Preference = 200 : 1
```

With these preference values, the NAS identifies the attribute set tagged 2 as the primary attribute set, and first attempts to establish an L2F tunnel to a server at 2.2.2.2. It tries an L2TP tunnel to the LNS at 1.1.1.1 only if the initial tunnel attempt fails. If the second attempt also fails, the system attempts to establish an L2TP tunnel to 3.3.3.3.

RADIUS: Support for MS-CHAP authentication

In previous releases, MS-CHAP authentication was supported only in local Connection profiles, not in RADIUS profiles. In addition, the MAX TNT did not supply a key for encrypting passwords by means of Data Encryption Standard (DES), so MS-CHAP did not work with older Microsoft software, such as LAN Manager and Windows 95.

With MAX TNT TAOS 8.0.0, the MAX TNT provides a key for DES encryption of passwords when MS-CHAP authentication is in use. For a description of how the key is used, see RFC 2433, *Microsoft PPP CHAP Extensions*. No new parameters are required in local profiles. For details about configuring MS-CHAP in local profiles, see the *MAX TNT Network Configuration Guide*.

For RADIUS profiles, the MAX TNT now supports the MS-CHAP-Challenge and MS-CHAP-Response vendor-specific attributes (VSAs) when requesting RADIUS MS-CHAP authentication. RFC 2548, *Microsoft Vendor-specific RADIUS Attributes*, defines the format and vendor-specific values for these attributes.

RADIUS: Authentication of Telnet sessions

You can now use the following attribute-value pair to specify a default User profile for RADIUS-authenticated Telnet access to the MAX TNT unit:

RADIUS attribute	Value
Ascend-Telnet-Profile (91)	Name of a MAX TNT User profile to be used for authenticating Telnet logins.

When a user attempts to Telnet into the MAX TNT interface, the system first looks for a User profile matching the login name and password given by the user. If that fails, the system uses the server specified in the External-Auth profile to locate a RADIUS user profile. If the RADIUS server returns a profile that includes the Ascend-Telnet-Profile attribute, the system uses the specified User profile to authenticate and set permissions for the session. Only RADIUS profiles that specify this attribute can be used to authenticate a Telnet login to the MAX TNT interface. Following is a sample RADIUS profile that enables Telnet access to the MAX TNT with administrator permissions:

```
admin Password = "secret-pw"  
    Service-Type = Framed-User,  
    Ascend-Telnet-Profile = admin
```

RADIUS: Inclusion of data and transmit rates in access request

The Ascend-Data-Rate and Ascend-Xmit-Rate RADIUS attributes are now part of an Access Request packet. The information they contain is sent only if you do not authenticate with CLID or DNIS. The Ascend-Data-Rate attribute specifies the receive rate of the connection in bits per second. The Ascend-Xmit-Rate attribute specifies the transmit rate for the connection. Following are the modified attributes.

Ascend-Data-Rate (197)

Description: The Ascend-Data-Rate attribute specifies the receive rate of the connection in bits per second.

Usage: Ascend-Data-Rate does not appear in a user profile. Its default value is 0 (zero). This attribute appears in Accounting-Request packets to provide troubleshooting information for the user.

Dependencies: The MAX TNT includes Ascend-Data-Rate in an Accounting-Request packet when both of the following conditions are true:

- The session has ended or has failed to authenticate because Acct-Status-Type is set to Stop.
- The Auth-Type parameter is not set to RADIUS/LOGOUT.

The MAX TNT includes Ascend-Data-Rate in an Access Request packet unless you authenticate with CLID or DNIS.

See Also: Ascend-Xmit-Rate

Ascend-Xmit-Rate (255)

Description: Specifies the rate of data transmitted on the connection in bits per second. For ISDN calls, Ascend-Xmit-Rate indicates the transmit data rate. For analog calls, it indicates the modem baud rate at the time of the initial connection.

Usage: Ascend-Xmit-Rate does not appear in a user profile. Its default value is 0 (zero). This attribute appears in Accounting-Request packets to provide troubleshooting information for the user.

Dependencies: The MAX TNT unit sends the Ascend-Xmit-Rate attribute in Accounting-Request packets at the end of a session under these conditions:

- The Accounting-Request packet has Acct-Status-Type set to Stop.
- The Auth-Type parameter is set to a value other than RADIUS/LOGOUT.

The MAX TNT unit sends the attribute with the Accounting-Request packet whether the unit authenticates the connection or not.

The MAX TNT unit includes Ascend-Xmit-Rate in an Access Request packet unless you authenticate with CLID/DNIS.

See Also: Ascend-Data-Rate

RADIUS: TCP-Clear-Quiet value added to Login-Service attribute

The new TCP-Clear-Quiet value to the Login-Service attribute enables the terminal-server software to suppress, on a per-user basis, status messages sent out to IP hosts upon establishment of a TCP-Clear connection.

By default, the terminal server returns a Connected status message when the MAX TNT establishes an incoming TCP-Clear connection. In previous releases, you could suppress all status messages for all login connections by enabling Silent-Mode in the Terminal-Server profile. Now, you can suppress status messages on a per-user basis for RADIUS-authenticated TCP-Clear connections.

Following is the modified RADIUS attribute.

Login-Service (15)

Description: The Login-Service attribute specifies the type of terminal-service connection established with an IP host immediately after authentication.

Usage: Assign one of the following values:

- Telnet (0), for a Telnet session with the host specified by the Login-Host attribute.
- Rlogin (1), for an Rlogin session with the host specified by the Login-Host attribute.
- TCP-Clear (2), for a TCP connection with the host specified by Login-Host, over which the user can run an application specified by Login-TCP-Port. If you specify this setting, the TCP-Clear-Answer Enabled value must be set to yes in the Answer-Defaults profile.
- TCP-Clear-Quiet (256), for a TCP-Clear session for which the terminal-server software suppresses status messages while the session is being established.

When you set the Login-Service attribute, a dial-in terminal-server user makes an immediate connection to an IP host on your local network and never sees the terminal-server interface.

By default, the MAX TNT does not grant immediate access to an IP host.

Dependencies: Keep in mind the following additional information:

Built-in features in MAX TNT TAOS 8.0.0

Authentication, Authorization, and Accounting (AAA)

- If you specify both Login-Service and Login-Host, the MAX TNT automatically connects the Login-User to the host specified by Login-Host.
- If you do not specify Login-Service or Login-Host, the Login-User sees either the MAX TNT unit's terminal-server command-line interface or the terminal-server menu interface, depending upon how you configure the MAX TNT.

Example:

Following is a sample profile that specifies TCP-Clear-Quiet login service.

```
tcpapp1 Password = "localpw"
  Service-Type = Login-User,
  Login-Service = TCP-Clear-Quiet,
  Login-Host = 10.10.10.1,
  Login-TCP-Port = 23,
  Login-Host = 10.10.10.2,
  Login-TCP-Port = 125
```

See Also: Login-Host (14), Login-TCP-Port (16)

RADIUS: Toggle boot requests

Sites that do not support a RADIUS boot server can now prevent the MAX TNT from sending unrequested management traffic in their networks. Following is the relevant parameter, shown with its default setting:

```
[in EXTERNAL-AUTH:rad-auth-client]
allow-auth-config-rqsts = yes
```

With this parameter's default setting (yes), the MAX TNT sends external configuration requests to the RADIUS server (the Boot Auth server). If you set this parameter to no, the MAX TNT does not send the external configuration requests.

RADIUS: Overriding the Answer-Defaults authentication method

With MAX TNT TAOS 8.0.0, you can use the Ascend-Auth-Type (81) vendor-specific attribute (VSA) in a RADIUS user profile to specify the type of PPP authentication to use, overriding the Answer-Defaults specification.

Some customers and providers who buy access from ISPs want to use CHAP authentication for their PPP calls, while other customers want to use PAP. In most cases, making both PAP and CHAP available to customers presents no problem. However, customers that use Microsoft Windows 95, Windows 98, or Windows NT clients cannot configure their units to reject CHAP. When Receive-Auth-Mode is set to any-ppp-auth, the MAX TNT offers CHAP authentication before PAP. Therefore, the Windows clients always use CHAP. In this release, you can configure RADIUS to select a different type of PPP authentication.

The Ascend-Auth-Type attribute is returned as part of the authorization resulting from DNIS or CLID first-tier authentication. If you specify a value for Ascend-Auth-Type, it overrides the Receive-Auth-Mode specification in the Answer-Defaults profile.

Ascend-Auth-Type (81)

Description: Specifies the type of PPP authentication the connection uses during first-tier CLID or DNIS authentication.

Usage: Specify one of the following settings:

- Auth-None (0) specifies that no second-tier name and password authentication is required. Specifying this value has the same effect as setting Ascend-Require-Auth to Not-Require-Auth.
- Auth-Default (1) specifies that the connection uses the Receive-Auth-Mode setting.
- Auth-Any (2) specifies that the connection must use PAP, CHAP or MS-CHAP.
- Auth-PAP (3) specifies that the connection must use PAP. The remote end sends its password in the clear. The password is not encrypted.
- Auth-CHAP (4) specifies that the connection must use CHAP. The remote end does not send its password in the clear. A Message Digest Algorithm 5 (MD5) digest calculated from the password and a random challenge are sent instead.
- Auth-MS-CHAP (5) specifies that the connection must use MS-CHAP.

If values other than those described above are passed from RADIUS to the MAX TNT, then the MAX TNT uses the Answer-Defaults profile default or the factory default.

ISDN callback

Callback is a feature in which unit A places a call to unit B, which hangs up and calls back unit A. The callback feature helps to make sure that the originating caller does not pay for the call and that the MAX TNT makes a connection with a known caller. Hanging up and calling back adds a level of certainty that the connection is with a trusted user, because the MAX TNT immediately calls back after verifying the user's name and password. To support callback, the MAX TNT must support both incoming and outgoing calls.

The MAX TNT supports the following three implementations of callback:

- CLID or DNIS callback (previously called Ascend CLID/DNIS Callback). The MAX TNT detects callback during the ringing state of an incoming call by means of the CLID or DNIS information element. The MAX TNT does not answer the call (go off hook), and the originating caller is not charged for the call.
- Ascend callback. This implementation is similar to CLID or DNIS Callback except that the MAX TNT detects callback during the authentication phase (after going off hook), by means of the username and password in the Connection profile. The originating caller is charged for the *initial* call.
- Callback Control Protocol (CBCP). This implementation was developed by Microsoft to address a need for greater security with PPP connections. The callback option defined in RFC 1570 is not as secure as other forms of callback because authentication is performed only during the initial call and *not* during the callback. CBCP callback, like Ascend callback, support a more secure connection, because the callback occurs *after* authentication.

CBCP offers features not available with the standard callback defined in RFC 1570. The client side supports a configurable time delay to allow users time to initialize modems or enable supportive software before the MAX TNT calls the client. The MAX TNT does not support a Connection profile or CBCP RADIUS profile to be shared by more than one Windows client.

The MAX TNT detects and negotiates CBCP callback by means of the CBCP protocol during PPP negotiation. You can configure the MAX TNT so that the user can negotiate the callback telephone number. CBCP callback takes the place of the Remote Access Server (RAS) server for callback to the RAS client (Windows 95).

Callback characteristics

Callback on the MAX TNT has the following characteristics:

- *Local or external authentication:* Connection profiles for configuring authentication can be either local or external (on the RADIUS server). For each implementation of callback, the MAX TNT accesses RADIUS once during external authentication. The MAX TNT does not request RADIUS attributes during the callback dialout process.
- *Nailed, Frame Relay or X.25 connections:* The MAX TNT does not support callback through nailed, Frame Relay, or X.25 connections.
- *Security involving external filters and routes:* The MAX TNT supports callback and external filters and external routes, but connections are restricted to a maximum of 16 external filters and 10 external routes per Connection profile.
- *ATMP tunnel security:* The MAX TNT supports Ascend Tunnel Management Protocol (ATMP) tunneling and all three callback implementations. The MAX TNT does not create the tunnel during the initial call, but creates the tunnel when the MAX TNT (foreign agent) calls back the mobile source.
- *Security:* You can use callback to extend security, and the MAX TNT clears all incoming calls that have callback enabled. If the MAX TNT cannot register the callback connection (for example, because of a lack of internal resources), the MAX TNT clears the call.
- *Expect callback:* This feature provides callback in reverse on the MAX TNT. The MAX TNT calls a remote unit, which calls back the MAX TNT. When the MAX TNT rejects the call, the MAX TNT disables its dialout process (to the dialed unit) until the unit calls back or 90 seconds is reached.
- *Callback log messages created:* Five new log messages provide information about callback processes. (See “Callback log messages” on page 85.)
- *Disconnect cause codes:* Incoming calls registered for callback are cleared with cause code 6 or 102. (Note: Unlike the MAX 4000, during the incoming call the MAX TNT does not log the message *LAN SESSION UP* when an Ascend callback is requested. Instead, the call is cleared with cause code 6, and the message *Callback-Registered* is logged.
- *Callback debug commands:* The MAX TNT supports new callback command to provide diagnostic information. (See “Callback debug command” on page 86.)
- *Special routing:* The MAX TNT can route a callback through a different resource type than the initial call. For example, the MAX TNT can accept the initial call from a modem card and the MAX TNT can make the outgoing call through a Hybrid Access card.

General information on configuring callback

Although you can mix callback implementations for a particular platform, you cannot mix callback types within the same connection profile. If you select more than one callback type, the MAX TNT performs callback in the following order: CLID or DNIS callback, Ascend callback, then CBCP callback. The MAX TNT supports CLID or DNIS authentication in combination with CBCP callback. Callback and expect-callback cannot be mixed inside the same profile, because you cannot simultaneously wait and perform a callback for a given profile.

For each callback implementation, an IP address pool index can be selected in place of a static IP address. The pool index enables the user to have a dynamic IP address taken from a pool. The MAX TNT assigns the IP address when the MAX TNT calls back the user.

Configuring CLID or DNIS callback

With CLID, a call comes in and the MAX TNT retrieves the matching profile with the Calling-Station-ID information from the ISDN Setup packet. The MAX TNT terminates the incoming call without answering it, and the MAX TNT initiates the callback. The devices negotiate PPP. With DNIS, the MAX TNT retrieves the matching profile with the Called-Station-ID. The originating caller cannot detect that the reason for the call termination is the pending callback, unless you have enabled the expect-callback feature. The expect-callback feature allows you to delay the originating caller from re-dialing the connection for ninety seconds.

Global parameter configuration

You must configure global configuration for CLID or DNIS callback within the Answer-Defaults profile:

Answer-Defaults parameter	Required settings
CLID-Auth-Mode	CLID-Require or DNIS-Require are required settings.

Local Connection profile configuration

You can configure local configuration through the Connection profile. Following are typical (and mandatory if indicated) Connection profile settings for CLID or DNIS callback, for a particular user:

Connection profile parameter	Typical setting
Active	Yes
Encapsulation-Protocol	PPP. Alternatively, you can select MP or MPP.
Dial-Number	The number to be dialed during the callback dialout phase.
CLID	The CLID number. To support CLID callback, you must specify a valid value for CLID.
Telco-Options>Callback	Yes.
Telco-Options>Data-Service	For example, Modem.
Telco-Options>Dialout-Allowed	Yes.
Telco-Options>Delay-Callback	For example, 10. This setting indicates the number of seconds that must elapse before the MAX TNT calls back the user.
CalledNumber	The DNIS number. To support DNIS callback, you must specify a valid value for CalledNumber.

External Connection profile configuration

Following are typical (and required if indicated) RADIUS Connection profile settings for CLID or DNIS callback.

RADIUS attribute	Typical setting
Password (2)	The password for CLID number. For example, Ascend-CLID.

RADIUS attribute	Typical setting
User-Service (6)	Dialout-Framed-User (5).
Ascend-Require-Auth (201)	Require-Auth (1).
Ascend-Callback (246)	Callback-Yes (1).
Caller-Id (31)	Specifies the calling-party number for Calling-Line ID (CLID) authentication. Indicates the telephone number of the user that wants to connect with the MAX TNT.
Framed-Protocol (7)	PPP (1), MP, or MPP (256).
Framed-Address (8)	The IP address of a caller. RADIUS can authenticate an incoming caller by matching the user's IP address to the one specified in the user profile. For example, 192.168.143.2
Framed-Netmask (9)	A subnet mask for the caller at Framed-Address. For example, 255.255.255.255
Ascend-Dial-Number (227)	The telephone number that the MAX TNT dials.
Ascend-Data-Svc (247)	Switched-Modem (42). Specifies the type of data service that the link uses for outgoing calls.
Ascend-Send-Auth (231)	Send-Auth-PAP (1)
Ascend-Send-Passwd (232)	The password that the RADIUS server sends to the remote end of a connection on an outgoing call. For example, Ascend.
Ascend-Route-IP (228)	Route-IP-Yes (1) or Route-IP-No (0). When you set this attribute to Route-IP-Yes (the default), IP routing is enabled for the profile. When set to Route-IP-No, IP routing is disabled for the profile.

Because the MAX TNT makes only one RADIUS request, all parameters must be present in the profile.

Expect-callback configuration

With CLID or DNIS callback, the MAX TNT hangs up on an incoming caller and immediately initiates callback. Callback ensures that a connection is made with a known destination. For outgoing calls, the call originator can be configured to expect a callback from the machine that is called. The expect-callback feature prevents the call originator from dialing out more than one time before being called back.

For example, a call is initiated by a MAX TNT to a Pipeline unit. The Pipeline unit receives an incoming ISDN Setup message, recognizes the CLID, and rejects the incoming call with a Disconnect message. If the expect-callback feature is set on the MAX TNT, it waits ninety seconds for the Pipeline to call back. If the feature is not set, the MAX TNT might determine that the call never got through and redial the call immediately.

When you set Expect-Callback to Yes on the calling device, all dialout calls that do not connect for any reason are put on a list that disallows further calls to that destination for 90 seconds. This delay gives the called device an opportunity to complete the callback.

To configure CLID or DNIS callback for expect-callback, two local connection parameters require configuration:

Connection profile parameter	Setting
Telco-Options>Callback	No
Telco-Options>Expect-Callback	Yes

Following are typical external Connection profile settings for expect-callback:

RADIUS attribute	Typical setting
Password (2)	Ascend
User-Service (6)	Dialout-Framed-User (5).
Ascend-Dial-Number (227)	The telephone number that the MAX TNT dials.
Framed-Protocol (7)	PPP (1).
Ascend-Data-Svc (247)	Switched-64K (2). Specifies the type of data service that the link uses for outgoing calls.
Ascend-Dialout-Allowed (131)	Dialout-Allowed (1).
Framed-Address (8)	The IP address of a caller. RADIUS can authenticate an incoming caller by matching the user's IP address to the one specified in the user profile. For example, 4.5.6.7.
Framed-Netmask (9)	A subnet mask for the caller at Framed-Address. For example, 255.255.255.0.
Ascend-Metric (225)	2. An integer that specifies the virtual hop count of an IP route. The default setting is 7.
Ascend-Send-Auth (231)	Send-Auth-PAP (1).
Ascend-Send-Passwd (232)	The password that the RADIUS server sends to the remote end of a connection on an outgoing call. For example, Ascend.
Ascend-Expect-Callback (149)	Expect-Callback-Yes (1).
Ascend-Route-IP (228)	Route-IP-Yes (1) or Route-IP-No (0). When you set this attribute to Route-IP-Yes (the default), IP routing is enabled for the profile. When set to Route-IP-No, IP routing is disabled for the profile.

Configuring Ascend callback

The MAX TNT performs Ascend callback after fully negotiating the PPP connection. When a MAX TNT is the called device, the call comes in and normal authentication occurs. The MAX TNT then terminates the call and initiates callback, the call is terminated, and the call negotiates PPP. When a MAX TNT is the calling device, it waits for a callback if the connection proceeds normally and is disconnected before any data passes. The MAX TNT does not redial for a specified number of seconds.

Global parameter configuration

Configure global configuration for Ascend callback within the Answer-Defaults profile. Following are typical global configuration settings for Ascend callback:

Built-in features in MAX TNT TAOS 8.0.0

Authentication, Authorization, and Accounting (AAA)

Answer-Defaults parameter	Typical setting
CLID-Auth-Mode	Cannot be set to CLID-Require or DNIS-Require. If the parameter is set to either of these, CLID or DNIS callback is performed instead of Ascend callback.
PPP-Answer>Enable	Yes.
PPP-Answer>Receive-Auth-Mode	Any-PPP-Auth.
PPP-Answer>CBCP-Enable	This setting is not required for Ascend callback.

Local Connection profile configuration

You configure local configuration within the Connection profile. Following are typical Connection profile settings for Ascend callback:

Connection profile parameter	Typical setting
Active	Yes.
Encapsulation-Protocol	PPP. Encapsulation-Protocol can also be set to MP or MPP.
Dial-Number	This parameter represents the number to be dialed during the callback dialout phase.
Telco-Options>Callback	Yes. If you specify CBCP-Enable, Telco Options > Callback takes precedence.
Telco-Options>Data-Service	Modem.
Telco-Options>Dialout-Allowed	Yes.
Telco-Options>Delay-Callback	For example, 10. This setting indicates the number of seconds that must elapse before the MAX TNT calls back the user.
PPP-Options>Send-Auth-Mode	Pap-PPP-Auth. Other PPP authentication (or None) can be used, depending on the remote side.
PPP-Options>Send-Password	For example, Ascend.
PPP-Options>Recv-Password	For example, Ascend.
PPP-Options>CBCP-Enabled	This parameter is not required for Ascend Callback.

External Connection profile configuration

Following are typical RADIUS configurations for Ascend Callback. For the displayed configurations, the use of external filters is assumed.

RADIUS attribute	Typical setting
Password (2)	The user password. For example, Ascend.
User-Service (6)	Framed-User (2).
Ascend-Callback (246)	Callback-Yes (1).
Framed-Protocol (7)	PPP (1), MP, or MPP (256).
Framed-Address (8)	The IP address of a caller. RADIUS can authenticate an incoming caller by matching the user's IP address to the one specified in the user profile. For example, 4.5.6.7.

RADIUS attribute	Typical setting
Framed-Netmask (9)	A subnet mask for the caller at Framed-Address. For example, 255.255.255.255.
Ascend-Dial-Number (227)	The telephone number that the MAX TNT dials.
Ascend-Data-Svc (247)	Switched-64K (2). Specifies the type of data service that the link uses for outgoing calls.
Ascend-Send-Auth (231)	For example, Send-Auth-PAP (1). An optional setting.
Ascend-Send-Passwd (232)	For example, Ascend. An optional setting. The password that the RADIUS server sends to the remote end of a connection on an outgoing call. If the value does not match the remote end's value (in the Connection>PPP Options>Recv-Password, or RADIUS user profile), the remote system rejects the call.
Ascend-Data-Filter (242)	<p>An optional setting. The Ascend-Data-Filter parameter specifies the characteristics of a data filter in a RADIUS user profile. The MAX TNT uses the filter only when it places or receives a call associated with the profile that includes the filter definition.</p> <p>IP Out Forward. This optional setting specifies an IP filter for filtering packets going out of the MAX TNT and specifies that the MAX TNT must forward each a packet that matches the filter.</p> <p>Generic Out Forward 12 ffff 0806. This optional setting specifies a generic filter for filtering packets going out of the MAX TNT and specifies that the MAX TNT must forward each packet that matches the filter. The offset, mask, and value are next specified.</p> <p>Generic Out Drop 0 0 0. This optional setting specifies a generic filter for filtering packets going out of the MAX TNT and specifies that the MAX TNT must drop each packet that matches the filter. The offset, mask, and value are next specified.</p>
Ascend-Route-IP (228)	Route-IP-Yes (1) or Route-IP-No (0). When you set this attribute to Route-IP-Yes (the default), IP routing is enabled for the profile. When set to Route-IP-No, IP routing is disabled for the profile.

Configuring CBCP callback

CBCP is an option negotiated during the Link Control Protocol (LCP) phase of PPP negotiation. Although you configure support for CBCP system wide on the MAX TNT, not every connection must negotiate CBCP callback. Parameters exist in the Answer-Defaults profile and in each Connection profile. The calling and called sides of a PPP session initiate authentication after acknowledging that CBCP is to be used.

The MAX TNT uses the username and password to link a caller with a specific Connection profile or RADIUS User profile. Configured CBCP parameters in the Connection profile specify variables for the callback. If, at any point, the client and the MAX TNT disagree about any CBCP parameters, the MAX TNT can drop the connection. The MAX TNT does not

support sharing a Connection profile or CBCP RADIUS profile by more than one Windows client.

Depending on the configuration, either the client or the MAX TNT can supply the callback telephone number.

In CBCP callback, a caller connects to the MAX TNT. LCP negotiations begin. The MAX TNT verifies that CBCP Mode is set in the profile. If the caller and MAX TNT successfully negotiate the LCP option for CBCP, CBCP then begins after authentication. The caller authenticates itself to the MAX TNT. If authentication fails, the MAX TNT terminates the connection. During CBCP, the client also supplies to the MAX TNT the number of seconds (a configurable value) it must wait before initiating the callback and, if applicable, the telephone number. The MAX TNT delays the callback on the basis of the previous negotiation. The MAX TNT dials the client by applying the information from the same profile used during negotiation.

Global parameter configuration

Configure global parameters within the Answer-Defaults profile. A new Answer-Defaults parameter, and two analogous RADIUS attributes, Ascend-CBCP-Enable and Ascend-CBCP-Mode, support CBCP callback. The CBCP-Enable parameter enables the CBCP protocol for incoming PPP calls.

Following are typical global configuration settings for CBCP Callback:

Answer-Defaults parameter	Typical setting
CLID-Auth-Mode	CLID-Require or DNIS-Require are not mandatory settings.
PPP-Answer>Enable	Yes.
PPP-Answer>Receive-Auth-Mode	Any-PPP-Auth.
PPP-Answer>CBCP-Enable	Yes is mandatory.

Local Connection profile configuration

You specify local configuration within the Connection profile. Following are examples of typical Connection profile settings for CBCP callback, for a particular user:

Connection profile parameter	Configuration
Active	Yes.
Encapsulation-Protocol	PPP. Encapsulation-Protocol can also be set to MP or MPP.
Dial-Number	If CBCP mode is set to CBCP-User-Number or CBCP-All, the callback telephone number can be given during callback negotiation. This setting can be left empty.
Telco-Options>Callback	No.
Telco-Options>Data-Service	For example, Modem.
Telco-Options>Dialout-Allowed	Yes.
PPP-Options>Send-Auth-Mode	Not required. Used with Windows 95, Windows 98, Windows NT.

Connection profile parameter	Configuration
PPP-Options>Recv-Password	For example, Ascend.
PPP-Options>CBCP-Enabled	Yes.
PPP-Options>Trunk-Group-Callback-Control	For example, 9. If the caller supplies the telephone number, set this parameter to the value that the MAX TNT prepends to the number supplied by the user when calling back.
PPP-Options>Mode-Callback-Control	CBCP-User-Number. This parameter has the following possible values: CBCP-No-Callback, CBCP-User-Number, CBCP-Profile-Num, and CBCP-All.

External Connection profile configuration

Following are typical RADIUS configurations for CBCP callback:

RADIUS attribute	Typical setting
Password (2)	The password for the CBCP user. For example, Ascend.
User-Service (6)	Framed-User (2)
Framed-Protocol (7)	PPP (1), MP (2) or MPP (256).
Ascend-Dial-Number (227)	The telephone number that the MAX TNT uses to call back when CBCP Mode is set to CBCP-Profile-Num or CBCP-All.
Ascend-Data-Svc (247)	Usually Switched-Modem (42), for CBCP. Specifies the type of data service the link uses for outgoing calls.
Ascend-Send-Auth (231)	For example, Send-Auth-None (0). Not a required setting.
Ascend-CBCP-Enable (112)	CBCP-Enabled (1).
Ascend-CBCP-Mode (113)	CBCP-Profile-Callback (3).
Ascend-Assign-IP-Pool (218)	1 (the default). An integer that corresponds to an address pool. When set to 0, RADIUS chooses an address from any pool that has one available. Set, if applicable.
Ascend-Route-IP (228)	Route-IP-Yes (1) or Route-IP-No (0). When you set this attribute to Route-IP-Yes (the default), IP routing is enabled for the profile. When set to Route-IP-No, IP routing is disabled for the profile.

The Ascend-CBCP-Trunk-Group setting is not mandatory. The setting is useful when the caller enters the callback number and trunk groups are used.

Callback log messages

Five new log messages have been implemented for the callback feature:

Log message	Description
Callback pending	Information message that occurs either when an incoming call arrives and is rejected because the user is already registered for callback, or when a non-callback outgoing call is requested and the user is registered for callback.
Callback registered	Information message that is generated when a user is registered for callback.
Callback dialout	Information message that occurs when the callback process performs the dial-out.
Callback register failed	Error message that occurs when the callback process is not able to register the callback.
Callback dialout failed	Error message that occurs when callback dialout fails.

Callback debug command

The Callback command provides debug information about callback operations. You can use the Callback command to debug the shelf controller or a host slot card. Because some options are card specific, each usage is documented separately.

Shelf controller card usage

```
admin> callback [-t|-l|-e|-?]
```

Syntax element	Description
-t	Toggle module debug level. Enables or disables the trace display.
-l	Display the external profile list. Displays the list of external Connection profiles (generally given by a RADIUS server) stored for the dial-out sequence. The lifetime of an external profile registered in the list is the same as the associated outgoing call duration.
-e	Display the callback list, which includes both callback and expect-callback entries. Displays the list of Connection profiles registered for callback.
-?	Display a summary of command options.

Host slot card usage

```
admin> callback [-t|-c|-f|-?]
```

Syntax element	Description
-t	Toggle module debug level. Enables or disables the trace display.
-c	Display the external profile cache. Displays the list of external Connection profiles (generally given by a RADIUS server) stored for the dial-out sequence. The list is managed as a data cache. The lifetime of an external profile registered in the list is limited to a few seconds.
-f	Flush the external profile cache. Allows the user to empty the cache.
-?	Display a summary of command options.

Management Agent

DOS-compatible FAT-16 flash memory format

In previous releases, the shelf controller PCMCIA flash memory cards used a proprietary format, which limited their use to the storage of code images. These cards could not store data, such as the system configuration. This release introduces a DOS-compatible general-purpose file system. In the initial release, the file system is supported on the MAX TNT shelf controller PCMCIA flash cards and Intel-compatible linear flash cards, but it has been designed with a minimum of platform dependencies.

The new flash format allows for hierarchical directories and eliminates the need to revise the file system format between versions. In addition, you can read and write the data on the flash card with a standard laptop or palmtop running OS/2 or a Windows version that supports Flash Translation Layer (FTL) linear flash memory.

File formats

The file allocation table-16 (FAT-16) file system is implemented on top of FTL. For details about the formats, see *PCMCIA Media Storage Formats, Chapter 5: Flash Translation Layer Microsoft FAT12 and FAT16 volume formats*.

Note: Filenames on MAX TNT flash cards must be compatible with the DOS 8.3 format.

A FAT-16 file system can store a large number of files in a hierarchy of directories. After you format flash under this software version, the flash card contains a top-level directory named `/current`, which contains the currently running version of the TAOS software as well as code image files for all supported slot cards. The slot card images are extracted from the tar file and stored as individual files with a `.ffs` filename extension. For example:

```
tntsr.ffs  
tnt8tl.ffs  
tnthdlc2.ffs
```

The new flash format also allows you to load a new software version or configuration data to the MAX TNT from a laptop running Windows or OS/2, rather than from a TFTP server. Because the FAT on FTL format is supported only on linear flash cards in this release, the laptop must have FTL linear flash.

Upgrading to the new flash file system

Because the MAX TNT TAOS 8.0.0 boot loader is FTL-aware, it can use either the older, proprietary flash format or the new FAT-16 format. When the FTL-aware boot loader is installed, you can switch to a previous software version (and back) by swapping out a flash card and resetting the system.

For flash cards that have been formatted to support the new FAT on FTL format, you can load a new software version onto a flash card from a laptop instead of a TFTP server.

To upgrade to MAX TNT TAOS 8.0.0 and use the new flash format, follow the instructions in “Upgrade and downgrade procedures” on page 23.

If you must downgrade to an earlier software version that does not support the FAT on FTL file system, follow the instructions in “Downgrade instructions” on page 28.

Creating directories in the flash file system

The `mkdir` command creates directories in the flash file system. The slash character (/) separates the elements of a pathname. For example, the following command creates a directory named `oldconf` at the top level of the flash card in slot 1:

```
admin> mkdir 1/oldconf
```

The following command creates a subdirectory named `conf1` within the `oldconf` directory:

```
admin> mkdir 1/oldconf/conf1
```

You can move files into a directory by using the `mv` command. For example, the following command moves a file named `0001conf` to the new subdirectory on flash card 1:

```
admin> mv 1/current/0001conf 1/oldconf/conf1/0001conf
```

Command changes for the new flash format

To support the new FAT-16 format, the following commands have been modified:

- Format
- Load
- Dircode
- Fack

Format command

The Format command formats flash memory in the new format by default. It supports a new `-o` option to format in the old proprietary format. Following is new command usage:

```
admin> format
prepare a flash card for use
usage: format [ -f ] [ -o ] < device >
[ -f ]:      (f)orce format without asking for verification
[ -o ]:      format in version 2("old") format
< device >: flash-card-1, flash-card-2
```

Load command

The Load command now supports an image type of `file` for Trivial File Transfer Protocol (TFTP) transfers to a flash card formatted for the FAT-16 format. Images of type `file` are not checked for an Internet Telnet Protocol (ITP) header, and are stored by name in the `/current` directory of the specified flash card. For example, the following command loads a voice-announcement file named `busy.au` from a TFTP server at 10.10.10.10 to the `/current` directory on flash card 1 (the default):

```
admin> load file network 10.10.10.10 busy.au
```

In addition, when used to load a tar file, the Load command now lists the filename of each code image in the file as the image is being extracted. For example:

```
admin> load tar network 10.10.10.10 tntrel.tar
file tntrel.tar...
untaring and loading image for...
shelf controller (tntsr/tntsr.ffs)...
```



```
8t1-card (tnt8t1/tnt8t1.ffs)...
skipping t3-card (tntt3/tntt3.ffs)...
4ether2-card (tntenet/tntenet.ffs)...
hdlc2-card (tnthdlc/tnthdlc.ffs)...
skipping 4swan-card (tntswan/tntswan.ffs)...
48modem-56k-card (tntmdm56k/tntmdm56k.ffs)...
skipping 48modem-card (tntmdm/tntmdm.ffs)...
skipping analog-modem2-card (tntamdm/tntamdm.ffs)...
done.
```

Dircode command

The Dircode command now notes the type and capacity of each card inserted. For an old format card, output include an indication of the type and capacity but is otherwise unchanged from output in previous versions. For example:

```
admin> dircode flash-card-1
Flash card code directory:
Card 1, format version 2, directory size 16, capacity 8MB
shelf controller reg    good  1231877 Dec 21 17:11      8.0.0
8t1-card reg           good   209191 Dec 21 17:11      8.0.0
4ether2-card reg       good   180385 Dec 21 17:11      8.0.0
hdlc2-card reg         good   588610 Dec 21 17:12      8.0.0
48modem-56k-card reg   good   724319 Dec 21 17:12      8.0.0
```

For a card with the new flash format, Dircode command output is slightly different from its output in previous versions, because information about each image is stored differently. For example:

```
admin> dircode flash-card-2
Card 2, format FTL/FAT, capacity 8MB
/current:
shelf controller      1231877 Tue Dec 21 17:17:22 1999 8.0.0
8t1-card              209191 Tue Dec 21 17:17:42 1999 8.0.0
4ether2-card          180385 Tue Dec 21 17:17:56 1999 8.0.0
hdlc2-card            588610 Tue Dec 21 17:18:38 1999 8.0.0
48modem-56k-card      724319 Tue Dec 21 17:19:58 1999 8.0.0
```

Fsck command

The Fsck command, when run on a flash card that has the new format, now prints a summary of file structures on the card. For example:

```
admin> fsck 2
Volume Stats:
Block Size: 512 (typical: 512)
Blocks Per Cluster: 3 (typical: 1, may be powers of 2 up to 16)
Reserved Blocks: 1 (typical: 1, but may be 0 - hundreds)
Number of FATs: 2 (must be 2)
Number of Root Directory Entries: 96 (typically between 32 and 224)
Total Blocks: 11264
Media Descriptor: f0 (ignored)
Volume Info calculated from values above:
Blocks Per Fat: 11
Fat Start Block: 1
Root Dir Start Block: 23
Data Start Block: 29
Number of Root Dir Blocks: 6
```

```
Number of Clusters: 3745
FAT Type: Fat12
Cluster Usage
Usable Clusters: 3743
Free Clusters: 1828
Clusters lost during interrupted writes: 0
Other reserved clusters: 1909
```

Additional onboard memory for extended profiling

One 128-KB segment of onboard flash memory that was previously reserved for the boot image can now be made available for storing configuration profiles. Extending the amount of memory used for profiles is referred to as *extended profiling*. When this feature is enabled, twice as many configured profiles can be stored in nonvolatile RAM (NVRAM).

Note: The system requires a 32-MB memory card to use extended profiling.



Warning: Before enabling extended profiling, you must complete the entire upgrade process so that both the boot loader and the shelf controller images support extended profiling. If you enable the feature while one of these images is at a version level that does not support extended profiling, *all profile information will be lost*.

Enabling extended profiling

To use extended profiling, follow these steps:

- 1 Verify that the system has a 32-MB memory card installed.
- 2 Upgrade to MAX TNT TAOS 8.0.0 (see “Upgrade and downgrade procedures” on page 23).
- 3 Use the NVRAM -u command to verify that the system can now support extended profiling. (A message indicating that extended profiles is not in use appears.)

```
admin> nvrाम -u
Not Using Extended profiles
NVRAM seg[0]:start 1000C028 size 131064 avail 117904 used 13160 cmpct 0
```

- 4 Use the NVRAM -e command to enable extended profiling:

```
admin> nvrाम -e
```

The system displays the following messages and a prompt:

```
Warning 1, Note this command will RESET THE SYSTEM.
```

```
Warning 2, if your tntsrб load does not support extended profiles,
you WILL Lose your profiles!
```

```
Warning 3, once you switch to extended profiles it is difficult to
switch back. Switching back from Extended Profiles will require
offloading the profiles, clearing the flash, resetting the system,
and reloading the profiles.
```

```
Warning 4, This feature also requires a 32M memory card.
```

```
Are you sure you want to switch to Extended Profiles? [y/n]
```

- 5 Respond to the prompt as appropriate. To leave the system unchanged, type **n**. To enable extended profiling (which causes a reset), type **y**.

```
Are you sure you want to switch to Extended Profiles? [y/n] y
```

```
Compacting NVRAM storage...
...compact segment 0 done
```

Please stand by. System reset in progress...

*** Ascend TNT, commencing second stage boot***

- 6 After the reset, use the NVRAM -u command to verify that the system is now using extended profiles.

```
admin> nvram -u
Using Extended profiles
NVRAM seg[0]:start 1002C028 size 131064 avail 118896 used 12168 cmpct 1
NVRAM seg[1]:start 1004C028 size 131064 avail 131064 used 0 cmpct 1
```

Two memory segments are now available for profile storage.

Downgrading to a version that does not support extended profiling

Once you have enabled extended profiling, you cannot restore the system's configuration to an earlier software version that does not support extended profiling if the configuration file is larger than one segment (128-KB).

Restoring an earlier configuration

Because software releases are not necessarily backward compatible, Lucent recommends that you restore a backup configuration made under the earlier software version or one of its predecessors. If you must restore the current configuration to an earlier software version, see "Restoring an extended profile configuration" on page 92.

To downgrade to an earlier software version after enabling extended profiling in MAX TNT TAOS 8.0.0, follow these steps:

- 1 Save the configuration to another system. For example:

```
admin> save network 10.10.10.10 cfgsave
```

- 2 Reformat NVRAM.

```
admin> nvram -f
```

Note: When you clear NVRAM, the system loses its IP addresses. A multishelf system also loses multishelf behavior. So, after reformatting NVRAM, you must perform the following steps to regain a minimal IP configuration and (if appropriate) multishelf capabilities:

- In the IP-Interface profiles, set the IP addresses of the system's Ethernet ports.
- In the System profile, set shelf controller-Type to Master.
- Reset the system for these changes to take effect.

- 3 Load the earlier version of the boot loader. For example:

```
admin> load boot-sr network 10.10.10.10 tntsr.b
```

Note: If you are downgrading a multishelf system, propagate the boot loader to the slave shelves by using the Loadslave command. (The version of the tntsr.b file on the master shelf must match the tntsr.b file version on the slave shelves. Otherwise, the slave shelves cannot load code from the master shelf.)

- 4 Load the earlier version of the tar file. For example:

```
admin> load tar network 10.10.10.10 tntrel.tar
```

- 5 Reset the system.

- 6 Restore a backup configuration made under the restored software version or one of its predecessors.

Restoring an extended profile configuration

If you must restore the system's current configuration after downgrading to an earlier software version, you must first decrease the total amount of memory used for profile storage to less than 128-KB. To do so, you might need to remove some profiles. You can also compact the profiles to save more space.

The following example shows a system using about half of the second NVRAM segment for profiles. To restore such a configuration to a system version that does not support extended profiling, you must remove or compact about 51 KB of profile storage, as shown in the following steps:

- 1 Use the `nvr` `-u` command to display the amount of memory used for profile information. For example:

```
admin> nvr -u
Using Extended profiles
NVRAM seg[0]:start 1008C068 size 131064 avail 768 used 130296 cmpct 0
NVRAM seg[1]:start 100AC068 size 131064 avail 79188 used 51876 cmpct 1
```
- 2 Start removing profiles that are not required. For example:

```
admin> del -f con junk11xxx
CONNECTION/junk11xxx deleted
admin> del -f con junk12xxx
CONNECTION/junk12xxx deleted
```
- 3 After deleting profiles, use the `nvr` `-c` command to compact memory storage, followed by another `nvr` `-u` command to see how much space has been saved. For example:

```
admin> nvr -c
Compacting NVRAM storage...
...compact segment 0 done
...compact segment 1 done
admin> nvr -u
Using Extended profiles
NVRAM seg[0]:start 1008C068 size 131064 avail 5586 used 125478 cmpct 1
NVRAM seg[1]:start 100AC068 size 131064 avail 127603 used 3461 cmpct 1
```
- 4 Save the revised configuration to another system. For example:

```
admin> save network 10.10.10.10 cfigsave
```
- 5 Reformat NVRAM.

```
admin> nvr -f
```

Note: When you clear NVRAM, the system loses its IP addresses. A multishelf system also loses multishelf behavior. So, after reformatting NVRAM, you must perform the following steps to regain a minimal IP configuration and (if appropriate) multishelf capabilities:

 - In the IP-Interface profiles, set the IP addresses of the system's Ethernet ports.
 - In the System profile, set shelf controller-Type to Master.
 - Reset the system for these changes to take effect.
- 6 Load the previous version of the boot loader. For example:

```
admin> load boot-sr network 10.10.10.10 tntsrbin
```

Note: If you are downgrading a multishelf system, propagate the boot loader to the slave shelves by using the Loadslave command. (The version of the `tntsrbin` file on the master shelf must match the `tntsrbin` version on the slave shelves. Otherwise, the slave shelves cannot load code from the master shelf.)

- 7 Load the previous version of the tar file. For example, to load via TFTP from a local host:

```
admin> load tar network 10.10.10.10 tntrel.tar
```

- 8 Reset the system.

- 9 Restore the configuration saved in Step 4.

```
admin> load config network 10.10.10.10 cfigsave
```

Pattern detection and reporting in the TCP-Clear data stream

You can now run the Tokencount diagnostic command to detect and report the number of instances of a specified pattern (a *token*) in the TCP-Clear data stream sent by the MAX TNT. On the shelf controller, the command enables or disables the token-counting process, specifies up to four patterns, clears counters, and displays token information system-wide. Updates to the command specified on the shelf controller are immediately propagated to the host cards.

Note: Running the token-counting process incurs a substantial system performance penalty. When token-counting is enabled, the system scans all outbound data sent to TCP-Clear sessions for a specified pattern, and increments a counter for each match. If the system resets, it loses the token information.

Tokencount command syntax

On the shelf controller, the Tokencount command supports the following syntax:

```
usage: tokencount -option [ params ]
-a          clear counter for (a)ll tokens
-c n        (c)lear counter for nth token
-d          (d)isable token counting in the TCP-CLEAR buffer
-e          (e)nable token counting in the TCP-CLEAR buffer
-i          display counter (i)nfo
-u n pattern (u)pdate type nth token pattern
-?          display this summary
```

Option	Description
-a	Set token counters to zero. If the system resets, all token counters are set to zero. If a card resets, counters on that card are set to zero.
-c n	Set the counter for the specified token to zero.
-d	Disable the token-counting process.
-e	Enable the token-counting process.
-i	Display the current token search information, including the number found of each defined token.
-u n	Define a search token pattern and assign it the specified number.

Each pattern can contain up to 20 characters, but the first specified character cannot be repeated in the pattern more than eight times. You can specify the pattern as a combination of

alphanumeric, hexadecimal, octal, and special characters, but output on the host is always in hexadecimal format. The following special characters are significant when specifying the pattern:

Characters	Meaning	ASCII value
\x##	Hex format	N/A. To insert a 2-digit hexadecimal number in the pattern, precede the number with \x.
\##	Octal format	N/A. To insert a 2-digit octal number, precede the number with a backslash.
\a	Alarm	7
\b	Backspace	8
\f	Form feed	12
\n	Newline	10
\r	Return	13
\t	Tab	9
\v	Vertical tab	11
\\	Backslash	92
\"	Quotation mark	34
\'	Apostrophe	44

Examples of using Tokencount

The following commands enable the token-counting process and define four token patterns:

```
admin> tokencount -e
admin> tokencount -u 1 \xB0\x35\xFF\x10\x01
admin> tokencount -u 2 LC\n
admin> tokencount -u 3 A1\12\15
admin> tokencount -u 4 \a\b\f\n\r\t\v\\\'\"
admin> tokencount -i
Tokencount is enabled
  Number of "\xB0\x35\xFF\x10\x01" token received:0
  Number of "LC\n" token received:0
  Number of "A1\12\15" token received:0
  Number of "\a\b\f\n\r\t\v\\\'\"" token received:0
```

The next commands open a session with a modem card in shelf 5, slot 6 and display the token information gathered on that card:

```
admin> open 5 6
csm3-5/6> tokencount
Tokencount is enabled
  "0xb00x350xff0x100x1" token received:0
  "0x4c0x430xa" token received:0
  "0x410x310xa0xd" token received:0
  "0x70x80xc0xa0xd0x90xb0x5c0x270x22" received:0
```

Tokencount error messages

When Tokencount is enabled, it can generate the following error messages:

error: token type index must be in the range of 1 to 4

The number specified in the Tokencount -u command is out of the valid range from 1 to 4.

error: max. token size is 20

More than 20 characters were specified as a pattern in the Tokencount -u command.

error: wrong token type index

The character immediately following Tokencount -u was not numeric.

Telnet access control list

To enable you to permit Telnet access to the MAX TNT only from specific IP addresses, the MAX TNT supports a new Telnet Access Control List (TACL) profile. You must have System authorization to create, read, or modify the profile.

You can configure up to 20 entries in the TACL profile. Each entry can specify a host address (with a /32 subnet mask) or a subnet address. Specifying a subnet address allows access from any of the addresses in the subnet range.

The TACL profile contains the following parameters, shown here with default values:

```
[in TACL]
enable-permit = no

[in TACL:permit-list[1]]
valid-entry = no
source-address = 0.0.0.0/0
source-address-mask = 0.0.0.0
```

Parameter	Specifies
Enable-Permit	Enable/disable control over Telnet access to the unit on the basis of the Permit-List settings in the TACL profile. If set to <code>no</code> (the default), the Permit-List settings have no effect. If set to <code>yes</code> , only the IP addresses specified in the Permit-Lists are allowed to telnet into the MAX TNT command-line interface. Setting Enable-Permit to <code>yes</code> has no effect if no Permit-Lists have been specified.
Valid-Entry	Enable/disable the Permit-List entry.
Source-Address	Source IP address of a host or subnet to be allowed Telnet access to the MAX TNT unit. The specified subnet mask determines whether the entry is valid for a single host or a subnet. If you specify the subnet mask as part of the Source-Address value, the Source-Address-Mask value is set automatically to the corresponding dotted decimal value.
Source-Address-Mask	The subnet mask to be applied to the Source-Address value before enabling a host Telnet access to the unit. You can set the value directly in dotted decimal format or by including a subnet as part of the Source-Address value.

For example, the following commands create a TACL profile that enables Telnet access from 30 host addresses from 10.27.34.1 to 10.27.34.31:

```
admin> new tacl
TACL read

admin> set enable-permit = yes
admin> set permit-list 1 valid-entry = yes
admin> set permit-list 1 source-address-mask = 10.27.34.1/27

admin> list permit-list 1
[in TACL:permit-list[1] (changed)]
valid-entry = yes
source-address = 10.27.34.1/27
source-address-mask = 255.255.255.224

admin> write
TACL written
```

Periodic log message for reporting the software version

To facilitate troubleshooting procedures, you can now configure the MAX TNT to log the current software version every hour, rather than at system startup only. Following is a sample log message:

```
LOG debug, Shelf 1, Controller, Time: 13:00:46--
Software version 8.0.0
```

Following is the relevant parameter, shown with its default value:

```
[in LOG]
log-software-version = no
```

Parameter	Specifies
Log-Software-Version	Enable/disable hourly log messages reporting the current software version. The message is sent to the Syslog host. If Debug permission is enabled, the message is also displayed on the screen.

Remote management of other units

The Remote command is now available in the terminal-server interface on host cards that accept digital calls, and as a command on the MAX TNT shelf controller. As on other TAOS platforms, the Remote command is used to remotely manage another unit.

Opening a remote management session

During a remote management session, the user interface of the remote device is displayed as if you had opened a Telnet connection to the device. For example:

```
admin> remote allwynp50
```


allwynp50 Edit		
Main Edit Menu Configure >00-000 System 20-000 Ethernet 30-000 Serial WAN	10-100 1	00-200 11:23:55
	Link A	M31 Line Ch
	B1 A	Outgoing Call
	B2	
	20-100 Sessions	20-500 DYN Stat
	>1 Active	Qual Good 01:23:44
		OK 1 channel
		CLU 100% ALU 100%
	20-300 WAN Stat	20-400 Ether Stat
	>Rx Pkt: 667435 ^	>Rx Pkt: 99871435
	Tx Pkt: 3276757	Tx Pkt: 76876757
	CRC: 323v	Col: 73298
	00-100 Sys Option	00-400 HW Config
	>Security Prof:1 ^	>BRI Interface
	Software +8.0+	Adrs: 00c05b45390
	S/N:4293801 v	Enet I/F: AUI

Press Ctrl-n to move cursor to the next menu item. Press return to select it.
Press Tab to move to another window--thick border indicates active window.

The Remote command argument is the station name, which must match the value of a Station parameter in a Connection profile, or the user ID at the start of a RADIUS profile. The connection must use the MP+ protocol, and the connection must already be established when you use the Remote command.

When you use the Remote command on the shelf controller, it locates the host card that has an active connection to the remote unit. It then opens a session to that card, invokes the terminal-server interface, and uses the Remote command on the card to bring up the remote management session. The Remote command uses a proprietary protocol to connect to the remote unit and bring up its LCD menu, which can be used to reconfigure the unit. However, because your initial permissions are set by the default Security profile on the remote system, you might need to authenticate the Full Access or other administrator-level Security profile before managing the unit.

You can also manually open a session with the host card that has an active connection to the remote unit, invoke the terminal-server, and run the Remote command on the slot card. For example:

```
admin> userstat -s
SessionID Line/Chan Slot:Item Tx/Rx Rate Svc Address Username
228687860 1.01.02/01 1:03:01/01 56K/56K PPP 10.100.0.1 clarap75
228687861 1.02.03/02 1:04:02/00 28800/33600 MPP 10.168.6.24 allwynp50

<end user list> 2 active user(s)

admin> open 1 4

hdlc2-1/4> terminal-server

ascend% remote allwynp50
```

Terminating a remote management session

To exit from the remote management session and return to the command-line interface session on the shelf controller, type Ctrl-C three times in quick succession.

If you opened the session on a slot card, press Ctrl-\ to end the session. You can then quit the terminal server and the slot card session to return to the shelf controller.

Either end of the connection can terminate an MP+ connection by hanging up all channels of the connection.

Note: A remote management session can time out, because the traffic it generates does not reset the idle timer. Therefore, the Idle parameter in the Connection profile at both the calling and answering ends of the connection must be disabled during a remote management session, and restored just before exiting. Remote management works best at higher terminal speeds.

Error messages

The MAX TNT generates an error message for any condition that causes the session to terminate before sending the full number of packets. The following error messages can appear:

Message	Explanation
not authorized	Permissions are insufficient for beginning a remote management session. You must authenticate a User profile that enables the System permission.
cannot find profile for <station>	No profile was found for the specified station name.
profile for <station> does not specify MPP	A profile was located for the specified station name, but it did not specify the MP+ encapsulation protocol.
cannot establish connection for <station>	The MP+ connection to the remote station could not be established.
<station> did not negotiate MPP	The remote station did not negotiate an MP+ connection. Possibly the profile for the MAX TNT dial-in did not specify MP+.
far end does not support remote management	The remote station is running a version of TAOS that does not support remote management.
management session failed	A temporary condition, such as premature termination of the connection, caused the management session to fail.
far end rejected session	The remote station was configured to reject remote management. (The Remote Mgmt parameter was set to no in the remote station's System profile.)

Command for displaying filter information

The Filterdisp command enables you to display information about filters in use for active sessions. The command uses the following syntax:

```
filterdisp usage: filterdisp <sessNum>
                  without <sessNum> : display all active sessions and
                                      their filter names
                  with <sessNum>    : display filter details of the session
```

Displaying filter information for all active sessions

With no arguments, the command output lists all active sessions with associated filter information. For example:

```
admin> filterdisp
```

```
ID      Username      Src Route-Filter Data-Filter Call-Filter TOS-Filter
-----
010    dialin-23    ext
016    dialin-4     ext
017    edleung      ext          < filters present >
018    jwebster     ext          < filters present >
019    pyan         loc          datfilt2      callfilt4    tostestfilt
020    guest        ext
021    pvc2         loc      route-pvc
022    pvc4         loc
023    pvc5         loc
<end user list> 9 active user(s)
```

The output displays a session ID number, username, and an indication of where the session was authenticated (local or external). Sessions authenticated by local profiles display the filter names specified in the Connection profile. Externally authenticated sessions, such as RADIUS sessions, have no associated filter names so they appear with a <filters present> notation. The columns in the command output provide the following information:

Output field	Specifies
ID	Identification number for the session.
Username	Name of the authenticated profile.
Src	Source of the profile: whether it is downloaded through RADIUS (ext) or is a local profile (loc).
Route-Filter	If a route filter has been applied to the session. For sessions authenticated locally, the name of the filter is supplied. For externally authenticated sessions, <filters present> indicates that a route filter has been applied. If blank, no route filter applies.
Data-Filter	If a data filter has been applied to the session. For sessions authenticated locally, the name of the filter is supplied. For externally authenticated sessions, <filters present> indicates that a data filter has been applied. If blank, no data filter applies.
Call-Filter	If a call filter has been applied to the session. For sessions authenticated locally, the name of the filter is supplied. For externally authenticated sessions, <filters present> indicates that a call filter has been applied. If blank, no call filter applies.
TOS-Filter	If a type of service (TOS) filter has been applied to the session. For sessions authenticated locally, the name of the filter is supplied. For externally authenticated sessions, <filters present> indicates that a TOS filter has been applied. If blank, no TOS filter applies.

Displaying filter details for a single active session

To display the filter details for a particular session, specify the session ID as an argument on the Filterdisp command line. (To obtain the session ID number, first use the Filterdisp command without an argument, as described in the preceding section.) If you specify an invalid session number, the command returns an error. For example:

```
admin> filterdisp 3
Error: Invalid user session ID
```

The following sample output shows that no filters are applied to the sessions:

```
admin> filterdisp 23
Hostname:      pvc5
No associated filters

admin> filterdisp 10
Hostname:      dialin-4
No associated external filters
```

In the following sample output, call filters have been applied to a session that was authenticated locally:

```
admin> filterdisp 22

Hostname:      pvc4
Call Filter
Direction: In

Forward = no
Type = Generic Filter
offset = 0
len = 0
more = no
comp-neq = no
dummyForPadding = 0
mask = 00:00:00:00:00:00:00:00:00:00:00:00
value = 00:00:00:00:00:00:00:00:00:00:00:00

Call Filter
Direction: Out

Forward = yes
Type = Generic Filter
offset = 0
len = 0
more = no
comp-neq = no
dummyForPadding = 0

mask = 00:00:00:00:00:00:00:00:00:00:00:00
value = 00:00:00:00:00:00:00:00:00:00:00:00
```

The following sample output shows filters applied to an externally authenticated session:

```
admin> filterdisp 17
Hostname:      edleung
searching for external filters...
Externally obtained filters exist

Data Filter
Direction: Out
```

```
Forward = yes
Type = IP Filter
protocol = 0
source-address-mask = 0.0.0.0
source-address = 0.0.0.0
destination-address-mask = 0.0.0.0
destination-address = 0.0.0.0
Src-Port-Cmp = none
source-port = 0
Dst-Port-Cmp = none
dest-port = 0
tcp-estab = no

Forward = yes
Type = Generic Filter
offset = 12
len = 2
more = no
comp-neq = no
dummyForPadding = 0
mask = ff:ff:00:00:00:00:00:00:00:00:00:00
value = 08:06:00:00:00:00:00:00:00:00:00:00
```

Customizable screen width up to 255 characters

The MAX TNT now allows command-line input and terminal-server banners up to 255 characters, rather than the previous limit of 80 characters. Horizontal scrolling of the command line allows viewing of commands and banners that are wider than the terminal display.

To set the width of the terminal display window for the current session, use the Screen command. To specify the width to use for every login to the command-line interface, use the Screen-Width parameter in a User profile.

Setting screen width for the current session

The Screen command enables you to specify the width of the screen. The command uses the following syntax:

```
screen -w <width>
```

The Width argument is a value from 80 to 256 and default is 80. For example:

```
admin> screen -w 256
```

The specified screen width is the number of characters that are visible without scrolling, including the system prompt and spaces following it. For example, if the screen width is 80 characters and the prompt is admin> (a 6-character prompt followed by a space), the maximum number of visible characters in a command is 72. If the user enters a long command, for example that has 100 characters, 28 of the characters will not be visible at any one time. The user can scroll to the characters not currently visible by moving the cursor left or right.

The following control sequence allows users to redraw the current line:

Control sequence	Effect
Ctrl-L, Ctrl-R	Redraw line

All existing control sequences continue to work as in previous releases. For details, see the *TAOS Command-Line Interface Guide*.

Customizing a User profile for screen width

To enable you to specify the screen width for all subsequent sessions, the following parameter (shown with its default setting) has been added to User profiles:

```
[in USER/""]  
screen-width = 80
```

Parameter	Specifies
Screen-Width	Number of characters allowed on a command line or terminal-server banner. An integer from 80 (the default) to 255.

Following is an example of how to customize a user's profile for a screen width of 120 characters:

```
admin> read user admin  
USER/admin read  
  
admin> set screen-width = 120  
  
admin> write -f  
USER/admin written
```

TCP-Clear login host IP address reported

The IP address of login hosts for TCP-Clear sessions is now reported in the Simple Network Management Protocol (SNMP) event Management Information Base (MIB), Syslog messages, and the output of the Userstat command.

SNMP event MIB changes

When a TCP-Clear connection is successfully established, the login host's IP address is specified in the eventUserIPAddress object in the SNMP callCleared event. The definition of the eventUserIPAddress object in the event MIB has been modified as follows:

```
eventUserIPAddress OBJECT-TYPE  
    SYNTAX      IpAddress  
    ACCESS      read-only  
    STATUS      mandatory  
    DESCRIPTION "IP address of the remote user or login host.  
    Applicable only if 'eventType' is serviceChanged(4)  
    nameChanged(5) or callCleared. Value of a TCP-Clear  
    login host IP address is returned once a TCP-Clear  
    connection was successfully connected earlier in  
    a serviceChanged event.  
    The value 0.0.0.0 is returned if address is unknown  
    or if not applicable."  
    ::= { eventEntry 13 }
```

Syslog messages

When a TCP-Clear session is terminated, the login host's IP address is displayed instead of the zero address (0.0.0.0) in the Syslog message. For example:

```
[3/7/2/0] STOP: 'johnfan'; cause 11.; progress 43.; host 10.1.1.1  
[MBID 2] [johnfan]
```

Userstat command output

For an active TCP-Clear session, the login host's IP address is displayed instead of the zero address (0.0.0.0) in the Userstat Address field. For example:

```
SessionID Line/Chan Slot:Item Tx/Rx Rate Svc Address Username  
286993415 3.01.08/012 3:07:03/000 26400/26400 TCP 10.1.1.1 johnfan  
  
<end user list> 1 active user(s)
```

Note: If the TCP-Clear connection fails (if the login attempt has not been successfully established between the MAX TNT and any of the specified login hosts), the Userstat command shows the zero address in the Address field.

Coredump generated by specific Warning messages

In previous releases, the system generated a coredump whenever a fatal error message or warning 101 through 121 occurred, provided that coredumps were enabled, a coredump server was defined, and the server was either running the ascendump utility or had it configured to run when a coredump request packet was received.

With MAX TNT TAOS 8.0.0, when coredumps are enabled and set up as described immediately above, you can specify an additional range of Warning messages that will cause a coredump. The following new parameters (shown with default values) enable you to specify an additional range of Warning message index values to cause a coredump:

```
[in DEBUG/{ any-shelf any-slot 0 }]  
min-warning-core-dump = 0  
max-warning-core-dump = 0
```

Parameter	Specifies
Min-Warning-Core-Dump	Minimum Warning message index value to cause a coredump. This value must be less than or equal to the Max-Warning-Core-Dump value. The default zero means that only Warnings from 101 to 121 cause a coredump. The valid range is from 1 to 9999.
Max-Warning-Core-Dump	Maximum Warning message index value to cause a coredump. This value must be greater than or equal to the Min-Warning-Core-Dump value. The default zero means that only Warnings from 101 to 121 cause a coredump. The valid range is from 1 to 9999.

For example, the following command specify that in addition to Warnings 101 through 121, Warnings 500 through 600 will generate a coredump:

```
admin> read debug { 1 1 1}  
DEBUG/{ shelf-1 slot-1 1 } read  
  
admin> set min-warning-core-dump = 500  
  
admin> set max-warning-core-dump = 600
```

```
admin> write
DEBUG/{ shelf-1 slot-1 1 } written
```

Changes to the Debug profile are effective immediately.

Event logging when an operator downs or resets a slot card

In previous releases, you could not determine whether a card was reset because it failed or because an operator reset it. With MAX TNT TAOS 8.0.0, the MAX TNT generates Syslog records and NVRAM records showing that an operator reset a card.

The MAX TNT generates new Syslog records when you use the following commands:

- Slot -b (reset a card)
- Slot -d (bring a card down)
- Slot -u (bring a card up)

When you use Slot -b or Slot -d, the MAX TNT generates new NVRAM records as well.

New Syslog records

The MAX TNT generates three new Syslog records with a level of Warning when you use particular options associated with the Slot command. No Syslog record is generated if you reset the MAX TNT by means of the Reset command.

Syslog message when operator resets a card

When you use the Slot -b command, the following Syslog record is generated:

```
LOG warning, Shelf shelf_num, Controller, Time: 14:27:20--
Slot shelf_num/slot_num bounced
```

Suppose you specify the card in slot 6 on shelf 1 by entering the following command:

```
admin> slot -b 1 6
```

The following Syslog record is generated:

```
LOG warning, Shelf 1 Controller, Time: 14:27:20--
Slot 1/6 bounced
```

Syslog message when operator brings down a card

When you use the Slot -d command, the following Syslog record is generated:

```
LOG warning, Shelf shelf_num, Controller, Time: 14:27:20--
Slot shelf_num/slot_num down
```

Suppose you specify the card in slot 6 on shelf 1 by entering the following command:

```
admin> slot -d 1 6
```

The following Syslog record is generated:

```
LOG warning, Shelf 1 Controller, Time: 14:27:20--
Slot 1/6 down
```


Syslog message when operator brings up a card

When you use the Slot -u command, the following Syslog record is generated:

```
LOG warning, Shelf shelf_num, Controller, Time: 14:27:20--  
Slot shelf_num/slot_num up
```

Suppose you specify the card in slot 6 on shelf 1 by entering the following command:

```
admin> slot -u 1 6
```

The following Syslog record is generated:

```
LOG warning, Shelf 1 Controller, Time: 14:27:20--  
Slot 1/6 up
```

New NVRAM records

The MAX TNT generates two new NVRAM records when you enter the Slot -b or Slot -d command. No NVRAM record is generated if a slot card is brought up by a Slot -u command, or if the MAX TNT is reset by means of the Reset command.

NVRAM log message when operator resets a card

When you use the Slot -b command, the following NVRAM record is generated:

```
SLOT CARD BOUNCED: Index: 220 Revision: 8.0 Slot shelf_num/slot_num  
Date: 04/22/1999.      Time: 12:35:05  
Card bounced by 10.40.40.94, user profile admin.
```

Suppose you specify the card in slot 6 on shelf 1 by entering the following command:

```
admin> slot -b 1 6
```

The following NVRAM record is generated:

```
SLOT CARD BOUNCED: Index: 220 Revision: 8.0 Slot 1/6  
Date: 04/22/1999.      Time: 12:35:05  
Card bounced by 10.40.40.94, user profile admin.
```

NVRAM log message when operator brings down a card

When you use the Slot -d command, the following NVRAM record is generated:

```
SLOT CARD DOWN: Index: 221 Revision: 8.0 Slot shelf_num/slot_num  
Date: 04/22/1999.      Time: 12:35:05  
Card downed by 10.40.40.94, user profile admin.
```

Suppose you specify the card in slot 6 on shelf 1 by entering the following command:

```
admin> slot -d 1 6
```

The following NVRAM record is generated:

```
SLOT CARD DOWN: Index: 221 Revision: 8.0 Slot 1/6  
Date: 04/22/1999.      Time: 12:36:54  
Card downed by 10.40.40.94, user profile admin.
```

Support for multiple Syslog servers

You can now configure up to three Syslog servers, and specify which severity level the records must have to be sent to a Syslog server.

The stream of records sent by a unit to a Syslog server is called a *Syslog stream*. A Syslog server is typically a UNIX machine running a Syslog daemon. The parameters controlling the Syslog stream are stored in the Log profile. In the past, you could direct Syslog streams to a single Syslog server, and the system would log all records. With MAX TNT TAOS 8.0.0, the MAX TNT can support up to three concurrent Syslog streams, and can filter each stream independently on the basis of a specified Syslog level. For example, you can now have one Syslog stream transfer all records, another one transfer records with a severity level of warning or above, and a third stream transfer records with a severity level of emergency. Each stream goes to a separate server.

Overview of Log profile settings

The Log profile contains a new Syslog-Level parameter and two Auxiliary-Syslog subprofiles. Each Syslog data stream is configured independently in the following manner:

- All the settings in the Log profile, except the Syslog-Format value, affect the first data stream. The Syslog-Format setting controls the format of all Syslog streams.
- The settings in the Auxiliary-Syslog [1] subprofile affect the second data stream.
- The settings in the Auxiliary-Syslog [2] subprofile affect the third data stream.

Following are the relevant parameters, shown with default values:

```
[in LOG]
syslog-level = info

[in LOG:auxiliary-syslog[1]]
syslog-enabled = no
syslog-level = info
host = 0.0.0.0
port = 514
facility = local0

[in LOG:auxiliary-syslog[2]]
syslog-enabled = no
syslog-level = info
host = 0.0.0.0
port = 514
facility = local0
```

The settings in the Auxiliary-Syslog subprofile affect an individual Syslog stream and override the values specified in the Log profile. For details about the Syslog-Enabled, Host, Port, and Facility parameters, see the *MAX TNT Reference Guide*.

Note: In the Log profile, the Save-Level parameter does not control the level of Syslog records sent to the server. Instead, it controls the level of records displayed in the log window when you use the Log command.

Parameter reference entries

Syslog-Level

Description: Indicates the level of log messages to direct to a specified Syslog server. Messages at or above the specified level are sent to the server.

Usage: Specify one of the following settings

Setting	Lowest-level message indicates
none	No log message is directed to the Syslog server.
emergency	The unit has an error condition and is unlikely to be operating normally.
alert	The unit has an error condition but is still operating normally.
critical	An interface has gone down or a security error has occurred.
error	An error event has occurred.
warning	An unusual event has occurred, but the unit is otherwise operating normally. For example, this type of message appears when a login attempt has failed because the user entered an incorrect username or password.
notice	Events of interest in normal operation have occurred (a link going up or down, for example).
info (the default)	State and status changes that are commonly not of general interest have occurred.
debug	Debugging information.

By default, Syslog records with a level of `debug` are filtered out, and records with a level of `Info` or above are transmitted to the Syslog server. If you set `Syslog-Level` to `notice`, messages with a level of `notice` or higher are sent to the Syslog server.

Example: `set syslog-level = notice`

Dependencies: Consider the following:

- The Syslog-Level value in the Log profile affects the first data stream.
- The Syslog-Level value in the Auxiliary-Syslog [1] subprofile affects the second data stream.
- The Syslog-Level value in the Auxiliary-Syslog [2] subprofile affects the third data stream.

Location: Log, Log > Auxiliary-Syslog

See Also: Syslog-Enabled, Syslog-Level, Host, Port, Facility

Full backup queue type now identified in Syslog

With MAX TNT TAOS 8.0.0, the Syslog warning regarding a full backoff queue now specifies whether the records affected are RADIUS accounting records or call-logging records.

RADIUS accounting and call logging each make use of a backoff queue. When the backoff queue fills up, the MAX TNT begins to discard records. In past releases, the unit issued a Syslog warning that looked like the following:

```
Backoff Q full, discarding user XXX[YYY]
```

XXX is the username and YYY is the session ID.

Because the warning had the same format for both RADIUS accounting and call logging, you had no way of identifying which queue was full.

In this release, the following message appears if the backoff queue for RADIUS accounting records is full:

```
Backoff Q full, discarding user XXX [YYY]
```

If the backoff queue for call logging is full, the following message appears instead:

```
Call Log Backoff Q full, discarding user XXX [YYY]
```

Grep-like capability added to certain commands

With MAX TNT TAOS 8.0.0, you can filter the output of certain commands to display only the information matching a specified pattern. This new functionality operates in a similar way to piping the output of the command to `grep` in UNIX.

The number of commands that support the grep-like capability changes on a regular basis as the functionality is integrated into the system. Any command that can produce a large amount of output is a candidate for supporting this functionality. Following is a representative list of commands that currently support it:

```
arptable  
briChannels  
cadslLines  
callroute  
dadslLines  
dir  
ds3AtmLines  
filterdisp  
hdlc  
help  
if-admin  
ifmgr  
ipcache  
list  
modem  
netstat  
oc3AtmLines  
ospf  
swanLines  
tlchannels  
uds3Lines  
userstat  
vdslchannels
```

Using the grep feature

To search for a particular pattern in command output, use the following syntax:

command | **grep** [-i] [-v] *expression*

Option	Description
<i>command</i>	Command that supports the grep feature.
grep	Display only information that matches the <i>expression</i> pattern.
-i	Use pattern matching that is not case sensitive.
-v	Display only information that does <i>not</i> match the <i>expression</i> pattern.
<i>expression</i>	Expression to use for pattern matching.

For the *expression* argument, the grep feature supports the following regular expressions, wildcard characters, and patterns:

Regular expression	Description
\	Turns off any special meaning of the following character.
.	Matches any single character in the input string.
*	Matches zero or more occurrences of the previous character.
Single or double quotation marks	Enclose a pattern that contains spaces or other quotation marks.
^	Specifies the beginning of a line.
\$	Specifies the end of line.
	Specifies a logical OR.
[...]	Specifies any one of the characters in a range.
(...)	Groups expressions.

To search for a character that is a wildcard, you must precede it with the backslash (\) character, even if the wildcard character is within the boundaries of quotation marks.

The output data from the command is scanned line by line. If the pattern you specify is encountered in the line, that line is displayed. In addition, the number of lines found matching the pattern are counted and displayed at the end of the command. Note that the column headers and footers might be omitted from the display if they do not match the pattern. However, error messages are exempt from pattern matching.

If you use the grep feature with a command that does not support filtering, the system does not display an error. The command output is simply not filtered.

Examples of command output using the grep feature

Suppose the Userstat command displays the following lines without filtering:

```
291933498 1.01.06/14 1:16:05/029 64000/64000 MPP 38.13.173.193 LIsbits217
291933429 1.08.05/19 1:16:03/011 64000/64000 PPP 38.13.167.201 LIsbits26
291905815 1.01.02/13 1:11:03/000 34000/26400 PPP 38.13.104.19 ra5
```

The following command displays only the output lines that match the case-insensitive string `lisbits26`:

```
admin> userstat | grep -i lisbits26
291933429 1.08.05/19 1:16:03/011 64000/64000 MPP 38.13.167.201 LIIsbits26
<grep> Found 1 line(s) matching search criteria
```

The following command displays only the output lines that *don't* match the expression LIIsbits26:

```
admin> userstat | grep -v LIIsbits26
291933498 1.01.06/14 1:16:05/029 64000/64000 MPP 38.13.173.193 LIIsbits217
291905815 1.01.02/13 1:11:03/000 34000/26400 PPP 38.13.104.19 ra5
<grep> Found 2 line(s) matching search criteria
```

The following command displays only output lines that contain the number 64 plus any number of other characters followed by the string PPP:

```
admin> userstat | grep 64.*PPP
291933429 1.08.05/19 1:16:03/011 64000/64000 PPP 38.13.167.201 LIIsbits26
<grep> Found 1 line(s) matching search criteria
```

The following command displays only output lines that contain the string PPP followed by any four characters and the number 13:

```
admin> userstat | grep PPP....13
291933429 1.08.05/19 1:16:03/011 64000/64000 PPP 38.13.167.201 LIIsbits26
291905815 1.01.02/13 1:11:03/000 34000/26400 PPP 38.13.104.19 ra5
<grep> Found 2 line(s) matching search criteria
```

The following command displays only output lines that contain the string PPP followed by a space character, any character, and the number 13:

```
admin> userstat | grep "PPP 38.13"
291933429 1.08.05/19 1:16:03/011 64000/64000 PPP 38.13.167.201 LIIsbits26
291905815 1.01.02/13 1:11:03/000 34000/26400 PPP 38.13.104.19 ra5
<grep> Found 2 line(s) matching search criteria
```

The following command displays only output lines that contain the string LIIsbits217 or LIIsbits26:

```
admin> userstat | grep LIIsbits217|LIIsbits26
291933498 1.01.06/14 1:16:05/029 64000/64000 MPP 38.13.173.193 LIIsbits217
291933429 1.08.05/19 1:16:03/011 64000/64000 PPP 38.13.167.201 LIIsbits26
<grep> Found 2 line(s) matching search criteria
```

CLID display option for Userstat command

In the past, the Userstat command displayed the Dialed Number Information Service (DNIS) number of the user session, but not the caller ID (CLID). With MAX TNT TAOS 8.0.0, the CLID is added to the display for switched calls. To cause the system to display the CLID, the %f option has been added to the following parameter, shown with its default value:

```
[in SYSTEM]
userstat-format = "%i %l %s %r %d %a %u %c %t %n"
```

To instruct the system to display the telephone number from which the user is calling in the output of the userstat -l command, add the %f option to the parameter's value. For example, the following command specifies output that includes the Session ID (%i), IP address (%a), username (%u), and %f (calling number):

```
admin> read system
SYSTEM read

admin> set userstat-format = %i %a %u %f
```

```
admin> write
SYSTEM written
```

Following is an example of the resulting Userstat output. The caller's telephone number (if available) appears under the Calling# field of the command output. For example:

```
admin> userstat -l
SessionID  Address  Username  Calling#
287695661  10.1.2.1  ed-p130   1119855014
<end user list> 2 active user(s)
```

Options for displaying the system log

With MAX TNT TAOS 8.0.0, the Log command can display the system log on screen with paged output (the output is passed to a more function before display) and supports the grep-like capability of searching for particular strings. For details about using the grep-like search function, see “Grep-like capability added to certain commands” on page 108.

If you use the Log command without options, or with the top or bottom argument, it opens the Log status window, as in previous releases. With the -p, -r, or -t options, it displays the log at the command line. If you invoke the Log command from the boot loader, which does not support status windows, the command displays the log at the command line by default.

To use the Log command, you must have System permissions. The Log command supports the following syntax:

```
admin> help log
log usage: log [top | bottom | [-p -r -t] ]
```

Option	Description
-p	Print the contents of the system log to screen, with the most recent entry first.
-r	Print the contents of the system log in reverse order, with the oldest log entry first.
-t	Truncate the command output to the screen width. Many log entries are longer than the standard 80 characters of terminal output. This option truncates the output of the command to the screen width as defined by the current width set by the Screen command. For related information, see “Customizable screen width up to 255 characters” on page 101.

With the -p option, the Log command displays the system log with the most recent log entry first. For example, the following is sample output:

```
admin> log -p

Time      Date      Source                Level  Description
11:11:25  11/16/1999 shelf-1/controller    notice Slot 1/10, state UP 2
11:11:20  11/16/1999 shelf-1/slot-10      info   Software version 8.0.0
11:11:20  11/16/1999 shelf-1/slot-10      info   Card serial number 914694348
11:10:15  11/16/1999 shelf-1/controller    notice Slot 1/5, state UP 2
11:10:10  11/16/1999 shelf-1/slot-5        notice 100BaseT: Link down
11:10:10  11/16/1999 shelf-1/slot-5        notice iel-5-3: Link down
11:10:10  11/16/1999 shelf-1/slot-5        notice iel-5-2: Link down
11:10:10  11/16/1999 shelf-1/slot-5        notice iel-5-1: Link down
```

Option for displaying line status

With MAX TNT TAOS 8.0.0, the Line command can display line status on screen with paged output (the output is passed to a `more` function before display) and supports the grep-like capability of searching for particular strings. For details about using the grep-like search function, see “Grep-like capability added to certain commands” on page 108.

If you use the Line command without options, or with the `all`, `enabled`, `top`, or `bottom` arguments, it opens the Line status window, as in previous releases. With the `-p` option, the command displays the status information at the command line.

To use the Line command, you must have System permissions. The Line command supports the following syntax:

```
admin> help line
line usage: line [ [all | enabled ] [ top | bottom] ] | [ -p ]
```

Option	Description
-p	Print line status information to screen.

With the `-p` option, the Line command displays line status information directly to screen. For example, the following is sample output for T1 lines:

```
admin> line -p

Address Line State CARR LOOP DS0 Channel Status      Signaling Type
1/01/01 ACTIVE    --  LOOP ..... inband
1/01/02 RED ALARM LOC  --  ..... r1-inband
1/01/03 ACTIVE    --  --  ----- inband
1/01/04 RED ALARM --  --  ..... isdn-nfas
1/01/05 RED ALARM LOC  --  ..... inband
1/01/06 DISABLED  --  --  @@@@@@@@@ @@@@@@@@@ @@@@@@@@@ inband
1/01/07 DISABLED  --  --  @@@@@@@@@ @@@@@@@@@ @@@@@@@@@ inband
1/01/08 DISABLED  --  --  @@@@@@@@@ @@@@@@@@@ @@@@@@@@@ inband
```

Following is sample output for E1 lines:

```
admin> line -p

Address Line State CARR LOOP DS0 Channel Status      Signaling Type
1/14/01 ACTIVE    --  --  .----- s----- e1-indian-signa
1/14/02 RED ALARM LOC  --  ..... e1-dpnss-signal
1/14/03 ACTIVE    --  --  .----- s----- e1-indian-signa
1/14/04 DISABLED  --  --  @@@@@@@@@ @@@@@@@@@ @@@@@@@@@ @@@@@@@@@
1/14/05 DISABLED  --  --  @@@@@@@@@ @@@@@@@@@ @@@@@@@@@ @@@@@@@@@
1/14/06 DISABLED  --  --  @@@@@@@@@ @@@@@@@@@ @@@@@@@@@ @@@@@@@@@
1/14/07 DISABLED  --  --  @@@@@@@@@ @@@@@@@@@ @@@@@@@@@ @@@@@@@@@
1/14/08 DISABLED  --  --  @@@@@@@@@ @@@@@@@@@ @@@@@@@@@ @@@@@@@@@
```

The command displays the following line status information:

Output field	Description
Address	Shelf/Slot/Line number of the line. This information was displayed in the Line window in previous releases, and is documented in the <i>MAX TNT Reference Guide</i> .

Output field	Description
Line State	Status of the line. This information was displayed in the Line window in previous releases, and is documented in the <i>MAX TNT Reference Guide</i> . In addition, the LB line-state indicator has been added to indicate that an E1 line is looped back via the <code>fe-loop</code> command on the E1 card.
CARR	(Carrier). If the system detects a loss of carrier on a line, LOC is displayed. If the line sees carrier, it displays dashes (--).
LOOP	(Loopback status). If the line is locally looped, LOOP is displayed. Otherwise, the column contains dashes (--).
DS0 Channel Status	State of the individual DS0 lines. This information was displayed in the Line window in previous releases, and is documented in the <i>MAX TNT Reference Guide</i> .
Signaling Type	The type of signaling in use on the line. This information was displayed in the Line window in previous releases, and is documented in the <i>MAX TNT Reference Guide</i> .

Dynamic remote filters

With MAX TNT TAOS 8.0.0, you can create RADIUS pseudo-user profiles that define data filters. You can then apply the filters to multiple local Connection or RADIUS profiles by referring to the pseudo-user profile name.

When the MAX TNT receives a Filter-ID in an Access-Accept packet from RADIUS, it searches for a matching local filter. If it does not find one, the MAX TNT requests the filter from the RADIUS server. You can specify how the system should behave if the filter referred to in a profile is not found. The system can either establish the session and log a message about the missing filter, or terminate the call if a filter is not found.

Externally defined filters are cached locally for a configurable interval. The `FiltCache` command displays statistics about each cached RADIUS filter profile, and enables you to flush profiles from the cache. See “Command reference entry for remote filter cache management” on page 119 for information about the `FiltCache` command.

Current limitations

In this release, the remote filter implementation is subject to the following limitations:

- Filters applied to dialout calls are not supported in this release.
- Call filters, route filters, and TOS filters are not supported in this release. Only data filters are currently supported.

Overview of local profile settings

Following are the local parameters related to dynamic remote filters:

```
[in ANSWER-DEFAULTS:session-info]
filter-required = no

[in CONNECTION:session-options]
filter-required = no
data-filter = ""
```

```
[in IP-GLOBAL]
default-filter-cache-time = 1440
```

Parameter	Specifies
Filter-Required	Whether access to the filter is required for the session. With the default value of No, the system establishes the session even if the specified filter is not found. If the parameter is set to yes, the system disconnects the call if the filter is not found. This setting does not apply if the profile does not refer to a filter by name. The Answer-Defaults setting is used for RADIUS user profiles that apply a filter and do not explicitly specify a value for Ascend-Filter-Required (50).
Data-Filter	Name of a Filter profile associated with the connection. The name can be of a local profile or a filter pseudo-user profile in RADIUS. However, if a local Connection profile does not use authentication, it cannot specify a RADIUS filter profile.
Default-Filter-Cache-Time	Number of minutes to cache RADIUS filter profiles that do not include a value for Ascend-Cache-Time (57). The default is 1440 (24 hours). Once the cache timer expires, cached profiles are deleted from system memory. The next time a remote filter is needed, the system retrieves the profile from RADIUS and stores it in cache again. Keeping a profile in cache increases the performance of establishing sessions that use the filter, at the cost of some system memory. If this parameter is set to 0 (zero), the default timer is disabled so that only RADIUS profiles that specify a cache time are cached.

Overview of RADIUS user profile settings

RADIUS user profile support for filter profiles is provided by the following vendor-specific attributes (VSAs):

RADIUS attribute	Value
Filter-ID (11)	Name of a local or remote filter profile associated with the connection.
Ascend-Filter-Required (50)	Whether access to the filter is required for the session. With the default value of Required-No (0), the system establishes the session even if the specified filter is not found. If the attribute is set to Required-Yes (1), the system disconnects the call if the filter is not found. This setting does not apply if the profile does not refer to a filter by name. If this attribute is not specified, the Answer-Defaults setting is used to determine system behavior when the specified filter is not found.

Overview of RADIUS pseudo-user profile settings

A filter profile is a pseudo-user profile in which the first two lines have the following format:

```
profile-name Password = "ascend" Service-Type = Outbound
```

The *profile-name* value is any name you assign to the profile. Duplicate filter names are not allowed. If a local Filter profile is already stored, the MAX TNT does not retrieve a filter profile of the same name from the RADIUS server. Filter profile definitions can include the following attribute-value pairs:

RADIUS attribute	Value
Ascend-Data-Filter (242)	An abinary-format filter specification using one of the following formats: "generic <i>dir action offset mask value compare</i> [more]" "ip <i>dir action</i> [<i>dstip n.n.n.n/nn</i>] [<i>srcip n.n.n.n/nn</i>] [<i>proto</i>] [<i>destport cmp value</i>] [<i>srcport cmp value</i>] [<i>est</i>]]"
Ascend-Cache-Refresh (56)	Whether the timer for cached routes in this profile is reset each time a new session becomes active that refers to the pseudo-user profile. Refresh-No (0) does not reset the timer. Refresh-Yes (1) resets the cache timer when a session referring to the profile becomes active.
Ascend-Cache-Time (57)	Number of minutes to cache the profile. Once the cache timer expires for a RADIUS profile, the profile is deleted from system memory. The next time it is needed, the system retrieves it from RADIUS and stores it in cache again. Keeping a profile in cache increases the performance of route lookups, at the cost of some system memory. The minimum possible cache time is 0 minutes, which causes the system to retrieve the profile for every route lookup in the table. This setting is usually not desirable. If this attribute is not specified, the IP-Global setting is used.

To use these attributes, the RADIUS server must support vendor-specific attributes (VSAs) and the MAX TNT must be configured in VSA compatibility mode. Following are the relevant settings:

```
[in EXTERNAL-AUTH]
auth-type = radius

[in EXTERNAL-AUTH:rad-auth-client]
auth-radius-compat = vendor-specific
```

For details about these settings, see the *MAX TNT Reference Guide*. For details about defining data filters in RADIUS, see the *MAX TNT Network Configuration Guide*.

Examples of configuring a filter profile in RADIUS

Following is a sample RADIUS filter profile:

```
filter-c Password = "ascend", Service-Type = Outbound
  Ascend-Cache-Time = 20,
  Ascend-Cache-Refresh = Refresh-Yes,
  Ascend-Data-Filter = "ip out forward tcp dstip 10.1.1.3/16",
  Ascend-Data-Filter = "ip out drop"
```

The cache timer has been set to 20 minutes, and the timer is reset each time the filter is applied to a session.

The following commands configure a default cache time for RADIUS filter profiles:

```
admin> read ip-global
IP-GLOBAL read

admin> set default-filter-cache-time = 180

admin> write
IP-GLOBAL written
```

Following is a sample RADIUS filter profile that makes use of the default because a value for Ascend-Cache-Time (57) is not explicitly specified:

```
filter-e Password = "ascend", Service-Type = Outbound
  Ascend-Data-Filter = "ip out forward tcp dstip 10.2.2.2/28",
  Ascend-Data-Filter = "ip out drop"
```

Examples of applying remote filters

The following commands modify a Connection profile so that the session uses a remote filter and the system disconnects the call if the filter is not found:

```
admin> read connection p50-v2
CONNECTION/p50-v2 read

admin> set session-options data-filter = filter-c

admin> set session-options filter-required = yes

admin> write
CONNECTION/p50-v2 written
```

Following is a sample RADIUS profile that applies the same filter profile with the same requirements. This profile also specifies how the filters must be cached for this connection:

```
p50-v2 Password = "my-password", Service-Type = Framed
  Framed-Protocol = PPP,
  Framed-IP-Address = 10.1.1.1,
  Framed-IP-Netmask = 255.0.0.0,
  Filter-ID = "filter-c",
  Ascend-Filter-Required = Required-Yes
```

The following commands configure the system to reject incoming calls when the RADIUS user profile specifies a filter that is not found, and the user profile does not explicitly say what to do if the filter is not found:

```
admin> read answer-defaults
ANSWER-DEFAULTS read

admin> set session-info filter-required = yes

admin> write
ANSWER-DEFAULTS written
```

Following is a sample RADIUS profile that makes use of the default because a value for Ascend-Filter-Required (55) is not explicitly specified:

```
p50-v2 Password = "my-password", Service-Type = Framed
  Framed-Protocol = PPP,
  Framed-IP-Address = 10.1.1.1,
  Framed-IP-Netmask = 255.0.0.0,
  Filter-ID = "filter-c"
```

Managing remote filters

Filters defined in RADIUS pseudo-user profiles are accessible in the command-line interface as if they were local Filter profiles. For example, in the following listing, the profiles named `filter-a` and `filter-b` are local Filter profiles, and the profile named `filter-c` is a filter profile obtained from RADIUS:

```
admin> dir filter
  464  01/04/2000 19:01:49  filter-a
  470  01/04/2000 19:10:57  filter-b
 3901  01/04/2000 20:01:50  filter-c
```

You can read and list the contents of the remote filters in the usual way, as if they were local profiles. For example:

```
admin> read filter filter-c
FILTER/filter-c read (read-only)

admin> list
[in FILTER/filter-c]
filter-name* = filter-c
input-filters = [ { no no generic-filter { 0 0 no no +
output-filters = [ { yes no ip-filter { 0 0 no no +
```

Note: You cannot change RADIUS filter specifications from the command-line interface.

You can delete RADIUS filter profiles by using the `delete` command. For example:

```
admin> delete filter filter-c
Delete profile FILTER/filter-c? [y/n] y
FILTER/filter-c deleted
```

Parameter reference entries

Default-Filter-Cache-Time

Description: Specifies the default time (in minutes) during which the RADIUS filter profile remains locally cached on the MAX TNT.

Usage: Specify an integer. The default is 1440 minutes (24 hours). If you specify 0 (zero), the system does not cache the profile.

Example: `set default-filter-cache-time = 720`

Location: IP-Global

See Also: Filter-Required

Filter-Required

Description: Specifies whether the MAX TNT establishes a call if the filter profile applied in the caller's Connection profile cannot be found locally or in RADIUS.

Usage: Specify `yes` or `no`. The default is `no`.

- `yes` specifies that the MAX TNT does not establish a call if the filter profile applied in the caller's Connection profile cannot be found locally or in RADIUS.
- `no` specifies that the MAX TNT establishes a call if the filter profile applied in the caller's Connection profile cannot be found locally or in RADIUS.

Example: `set filter-required = yes`

Dependencies: Consider the following:

- If the call needs to be brought down, the cause code 425 results. If the call is allowed to come up, the system logs a notice-level message that the filter cannot be found.
- If the Ascend-Filter-Required attribute is missing in the RADIUS user profile, the MAX TNT uses the Filter-Required value in the Answer-Defaults profile.

Location: Answer-Defaults > Session-Info, Connection > Session-Options

See Also: Default-Filter-Cache-Time

RADIUS attribute reference entries

Ascend-Cache-Refresh (56)

Description: Specifies whether successive references to a cached filter profile reset its cache timer.

Usage: Specify one of the following values:

- Refresh-No (0) specifies that the cache timer is not reset.
- Refresh-Yes (1) specifies that the cache timer is reset.

See Also: Ascend-Cache-Time

Ascend-Cache-Time (57)

Description: Indicates the time (in minutes) during which the filter profile remains cached.

Usage: Specify an integer. If you do not specify the Ascend-Cache-Time attribute in a filter profile, the profile is cached for the amount of time specified by the Default-Prt-Cache-Time parameter in the IP-Global profile.

See Also: Ascend-Cache-Refresh

Ascend-Filter-Required (50)

Description: Specifies whether the MAX TNT establishes a call if the filter profile applied in the caller's RADIUS user profile cannot be found.

Usage: In a RADIUS user profile, specify one of the following values:

- Required-No (0) specifies that the MAX TNT establishes a call if the filter profile applied in the caller's RADIUS user profile cannot be found.
- Required-Yes (1) specifies that the MAX TNT does not establish a call if the filter profile applied in the caller's RADIUS user profile cannot be found.

Dependencies: Consider the following:

- If the call needs to be brought down, the cause code 425 results. If the call is allowed to come up, the system logs a notice-level message that the filter could not be found.
- If the Ascend-Filter-Required attribute is missing in the RADIUS user profile, the MAX TNT uses the Filter-Required value in the Answer-Defaults profile.

See Also: Filter-ID (11)

Command reference entry for remote filter cache management

The FiltCache command displays statistics about each cached RADIUS filter profile, and enables you to flush profiles from the cache.

FiltCache

Description: Displays the number of times a cached RADIUS filter profile was used, and enables you to flush all filter cache buffers.

Permission level: User

Usage: `filtcache -s [filtername] | -f [-f]`

Option	Description
<code>-s [filtername]</code>	If filtername is not specified, the command display statistics for all cached filters. If it is specified, the command displays statistics only for the specified filter.
<code>-f [-f]</code>	Flush all cached filters. The second <code>-f</code> flag specifies that all filters are flushed without waiting for confirmation.

Example: `filtcache -s myfilter`

Filter Name	Time Created	Exp After(min)	Use Cnt	Refresh Cache
myfilter	18:44:30	10	2	No

The command immediately above displays how many times a filter named `myfilter` has been used. The following command flushes all cached filters:

```
admin> filtcache -f
```

```
Flush all cached filter profiles? [y/n] y
All 3 cached RADIUS filter profiles flushed.
```

The following command displays how many times all cached RADIUS filters have been used:

```
admin> filtcache -s
```

Filter Name	Time Created	Exp After(min)	Use Cnt	Refresh Cache
myfilter	20:01:50	1440	3	Yes
filter-b	21:03:34	10	2	No
filter-c	21:10:32	8	14	Yes

MaxTap support

The MaxTap option logs user session data to a server. It is available only in countries where this capability is required by law for all remote access equipment.

Userstat options to display address and username

With MAX TNT TAOS 8.0.0, the Userstat command supports the following new options:

- `-a`, to take the IP address of a session as input and display the associated session details.
- `-u`, to take a username and display the associated session details.

- `-o`, to restrict the Userstat command output to specified fields.

Following is the new command usage statement:

```
admin> help userstat
userstat usage: userstat -options [ params ] [ -o [format] ]
command options:
  -s show users (default)
  -k <sessionID> kill a user session
  -a <ipAddress> show the session with matching <ipAddress>
  -u <username> show the session with matching <username>
  -l wide format (> 80 characters)
  -d dump, do not pass output through more
format values:
One or More of the following format characters
%i SessionID
%l Line/Chan
%s Slot:Item
%r Tx/Rx Rate
%d Type of Service
%a Address
%u Username
%c ConnTime
%t IdleTime
%n Dialed#
default : %i %l %s %r %d %a %u %c %t %n
```

Using the -o format specifier option

Use the `-o` option with one or more format specifiers to display only the fields of interest. For example, for an active session, the Userstat command shows the following details:

```
admin> userstat
SessionID Line/Chan Slot:Item Tx/Rx Rate Svc Address Username
288532030 1.01.01/012 1:03:01/002 56000/56000 PPP 1.1.1.238 net1
<end user list> 1 active user(s)
```

If you use the `-o` option and indicate the codes for SessionID and Line/Channel information, the command shows only the following details:

```
admin> userstat -o %i %l
SessionID Line/Chan
288532030 1.01.01/012
<end user list> 1 active user(s)
```

Using the -a and -u options

Use the `-a` option to display information related to a known IP address. It requires an IP address argument on the command line. For example:

```
admin> userstat -a 1.1.1.238
SessionID Line/Chan Slot:Item Tx/Rx Rate Svc Address Username
288532030 1.01.01/012 1:03:01/002 56000/56000 PPP 1.1.1.238 net1
<end user list> 1 active user(s)
```

To display only the relevant username, include the `-o` option as follows:

```
admin> userstat -a 1.1.1.238 -o %u
```



```
Username
net1
<end user list> 1 active user(s)
```

Use the `-u` option to display information related to a known username. It requires a user-name argument on the command line. For example:

```
admin> userstat -u net1

SessionID Line/Chan Slot:Item Tx/Rx Rate Svc Address Username
288532030 1.01.01/012 1:03:01/002 56000/56000 PPP 1.1.1.238 net1
<end user list> 1 active user(s)
```

To display only the user's IP address, include the `-o` option as follows:

```
admin> userstat -u net1 -o %a

Address
1.1.1.238
<end user list> 1 active user(s)
```

Selectable call-logging server

In previous releases, you could specify up to three call-logging servers, but the MAX TNT always directed its logging information to one of the servers. The system used the first configured server unless it was unavailable, in which case it used the second, or if that was unavailable, the third server. Once it started using a server, the system continued to do so until that server became unavailable.

With MAX TNT TAOS 8.0.0, you can control to which server the MAX TNT sends its logging information, provided that the Call-Logging profile is properly configured and enabled. Following are the relevant parameters, shown with default settings:

```
[in CALL-LOGGING]
call-log-server-index = host-1
```

Parameter	Specifies
Call-Log-Server-Index	Which of the configured <code>call-log-host-N</code> settings are used as the active call-logging server. Valid values are <code>host-1</code> (the default), <code>host-2</code> , and <code>host-3</code> . If the MAX TNT cannot authenticate the specified server, it attempts to use the next configured server.

To enable you to make this choice from an SNMP management station, the `callLoggingCurrentServerFlag` in the `callLoggingServerEntry`, which is in the Ascend call-logging MIB, is now a read-write variable. The variable can be set to 1 (active) or 2 (standby). Following is the new definition:

```
callLoggingCurrentServerFlag OBJECT-TYPE
    SYNTAX  INTEGER {
        active(1),
        standby(2)
    }
    ACCESS  read-write
    STATUS  mandatory
    DESCRIPTION "Value indicates whether this entry is the current
        Call Logging server or not. The standby(2) is not
```

```
set-able it is a value to report the standby status
of the Call Logging server."
::= { callLoggingServerEntry 2 }
```

Enhanced reporting of HDLC errors

The `hdlc-p` command on a Hybrid Access card now prints statistics every second, rather than summarizing the statistics in the output report. This provides a more detailed picture of error conditions. In addition, the total number of open channels is displayed. Following is a sample of the command's output:

```
hdlc2-1/4> hdlc -p
send rcv sndq rcvq d scr bufr crc long overrun inex abort txund nopen
0 0 0 400 1000 1800 0 0 0 0 0 0 0
0 0 0 400 1000 1800 0 0 0 0 0 0 0
0 0 0 400 1000 1800 0 0 0 0 0 0 0
0 0 0 400 1000 1800 0 0 0 0 0 0 0
0 0 0 400 1000 1800 0 0 0 0 0 0 0
```

The sample output displays the following fields:

Field	Description
send	Total number of packets sent in the previous second.
rcv	Total number of packets received in the previous second.
sndq	Total number of packets currently queued for transmission.
rcvq	Total number of packets currently queued for reception.
dscr	Total number of buffers for which there is an accounting. This is useful for detecting a buffer leak. Currently there is a total of 1800 buffers.
bufr	Total number of buffers for which there is an accounting. This is useful for detecting a buffer leak. Currently there is a total of 1800 buffers.
crc	Total number of packets received with crc errors in the previous second.
long	Total number of packets received in the previous second that were too long. Currently the maximum packet length is 2048 bytes.
ovrun	Total number of packets received in the previous second that could not be saved because there were not enough buffers.
inex	Total number of packets received in the previous second that were not a multiple of eight bits (after zero extraction).
abort	Total number of packets received in the previous second that were aborted by the reception of at least seven ones.
txund	Total number of packets transmitted in the previous second that were aborted because buffer chains were not ready in time. This status should always be zero since chained buffers are not used.
nopen	Total number of HDLC channels currently open. An HDLC channel corresponds to one or more TDM channels.

SNMP: Partial support for ATM MIB RFC 2515

With MAX TNT TAOS 8.0.0, the MAX TNT supports Get operations on the following tables of the ATM MIB described in RFC 2515, *Definitions of Managed Objects for ATM Management*:

- ATM Interface configuration table
- ATM Interface DS-3 Physical Layer Convergence Protocol (PLCP) table
- ATM Interface transmission convergence (TC) sublayer table
- ATM Interface virtual channel link (VCL) configuration table
- ATM Interface ATM Adaptation Layer 5 (AAL5) virtual channel connection (VCC) performance statistics table

Set operations are not yet supported. In addition, the following SNMPv2-related changes were made to `rfc2514.mib`:

- The definition of `atmMIB` and `atmMIBObjects` were moved here from `rfc2515.mib`.
- All the definitions were modified to SNMPv1 Structure of Management Information (SMI).
- The SNMPv1 entries `atmNoTrafficDescriptor`, `atmClpNoTaggingNoScr`, and `atmClpTaggingNoScr` are deprecated.

The following SNMPv2-related changes were made to `rfc2515.mib`:

- The definitions of `atmMIB` and `atmMIBObjects` were moved to `rfc2514.mib`.
- All MIB fields with Current Status were changed to Mandatory.
- MAX-Access syntax was changed to Access for all the fields.
- Fields with read-create access were changed to read-write.
- Set functions are not supported on the following parameters, so they have been changed from read-write to read-only:
 - `atmInterfaceConfEntry` parameters
 - `atmVclReceiveTrafficDescrIndex`,
`atmVclTransmitTrafficDescrIndex` and `atmVclAdminStatus` in `atmVclTable`
 - `atmVccAalType`, `atmVccAal5CpcsTransmitSduSize`,
`atmVccAal5CpcsReceiveSduSize`, `atmVccAal5EncapType`,
`atmVclRowStatus`, `atmVclCastType` and `atmVclConnKind` in `atmVclTable`
- Read-write permissions were changed to read-only permission in the following tables:
 - `atmTrafficDescrParamTable`
 - `atmVplTable`
- The `atmVpCrossConnectTable` and `atmVcCrossConnectTable` tables are not supported.

SNMP: Trap for dropped call-logging packets

A new Enterprise alarm-class trap (number 41) has been defined to signal SNMP management stations that call-logging packets are being dropped. Following is the relevant parameter, shown with its default value:

```
[in TRAP/""]  
call-log-dropped-pkt-enabled = yes
```

Parameter	Specifies
Call-Log-Dropped-Pkt-Enabled	Enable/disable sending a trap when a change in status is detected related to dropping call-logging packets. If enabled (the default), the system generates a trap when the value of the <code>callLoggingDroppedPacketCount</code> variable in the call-logging MIB is changed from 0 to 1 (which indicates that packets are being dropped) or from 1 to 0 (which indicates that packets are no longer being dropped). SNMP management stations can obtain the value of the variable at any time by using SNMP Get.

SNMP: Configuration of TFTP port for network management

The `flashOperationTftpPort` object has been added to Ascend's Flash MIB to configure the TFTP port setting for environments in which a network management station is running more than one management application, with a TFTP server local to each application.

The `flashOperationTftpPort` object is defined in the Flash MIB and used in the `load-config`, `save-config`, and `tftp-load` Flash MIB operations. The object's default setting is 69, which is the default port for TFTP operations. The object is defined as follows in the Flash MIB:

```
flashOperationTftpPort OBJECT-TYPE  
    SYNTAX      INTEGER  
    ACCESS      read-write  
    STATUS      mandatory  
    DESCRIPTION  
        "This object defines the port # to use on the remote system  
        when starting a TFTP operation using a flashOperationCommand. The  
        default port is 69/(tcp/udp) Trivial File Transfer."  
    ::= { flashOperation 8 }
```

The new Flash MIB has the following structure:

```
| -      1  flashDevice                flashGroup.1  
| | -     1  flashDevices              flashGroup.1.1  
| \_     2  flashDeviceTable          flashGroup.1.2  
|   \_   1  flashDeviceEntry          flashGroup.1.2.1  
|       | - 1  flashDeviceSocket      flashGroup.1.2.1.1  
|       | - 2  flashDeviceController  flashGroup.1.2.1.2  
|       | - 3  flashDeviceControllerSocket flashGroup.1.2.1.3  
|       | - 4  flashDeviceSize        flashGroup.1.2.1.4  
|       | - 5  flashDeviceUsed        flashGroup.1.2.1.5  
|       | - 6  flashDeviceState       flashGroup.1.2.1.6  
|       | - 7  flashDeviceMaster      flashGroup.1.2.1.7  
|       | - 8  flashDeviceFormatStatus flashGroup.1.2.1.8  
|       +- 9  flashDeviceDescription  flashGroup.1.2.1.9  
| -      2  flashFileTable            flashGroup.2
```

_	1	flashFileEntry	flashGroup.2.1
-	1	flashFileIndex	flashGroup.2.1.1
-	3	flashFileSocket	flashGroup.2.1.3
-	4	flashFileSize	flashGroup.2.1.4
-	5	flashFileStatus	flashGroup.2.1.5
-	6	flashFileName	flashGroup.2.1.6
-	7	flashFileChecksum	flashGroup.2.1.7
-	8	flashFileVersion	flashGroup.2.1.8
-	9	flashFileAccess	flashGroup.2.1.9
+-	10	flashFileDateTimeStamp	flashGroup.2.1.10
_	3	flashOperation	flashGroup.3
-	1	flashOperationStatus	flashGroup.3.1
-	2	flashOperationCommand	flashGroup.3.2
-	3	flashOperationHost	flashGroup.3.3
-	4	flashOperationDestFileName	flashGroup.3.4
-	5	flashOperationSrcFileName	flashGroup.3.5
-	7	flashOperationSocket	flashGroup.3.7
+-	8	flashOperationTftpPort	flashGroup.3.8

SNMP: New Config-Change trap

A new Config-Change trap is supported with an integer value of 30. The trap is issued whenever the system configuration is modified or a new software version is loaded. An SNMP management station can now receive a Trap (30) and a string containing the date, time, and information about the user that changed the configuration. On the MAX TNT, the trap has the following format:

Date, Time, "Configuration changed by user profile (YYY)."

YYY indicates the name of the User profile.

In addition, the sysConfigChange entry has been added to the Ascend MIB for the systemStatusGroup. Following is the Object ID for this string:

```
iso(1).org(3).dod(6).internet(1).private(4).enterprises(1).  
ascend(529).systemStatusGroup(9).sysConfigChange(14)
```

This member is an octet string and contains the date, time, and User or Security profile name and index. Following is the parameter for enabling the trap (shown with its default setting):

```
[in TRAP]  
config-change-enabled = yes
```

Config-Change-Enabled

Description: Enables or disables the configuration-change trap (Trap 30). The trap is enabled by default, which causes the system to issue the trap whenever the system configuration is modified or a new software version is loaded. If the parameter is set to no, the system does not issue the trap for those events.

Usage: Specify yes or no. The default is yes.

- yes enables the trap.
- no disables the trap.

Example: set config-change-enabled = yes

Location: Trap

SNMP: Trap for console state change now displays IP address

With MAX TNT TAOS 8.0.0, the SNMP agent on the MAX TNT sends the console's IP address in addition to the console index in the Console-State-Change trap. The Console-State-Change trap now carries the information displayed in the following example:

```
1999-07-02 12:07:26 eng-fast-4.ascend.com [192.168.25.4] enterprises.529:
Enterprise Specific Trap (12)Uptime:0:16:43
enterprises.529.8.2.1.1.2=2
enterprises.529.12.2.1.4.2=IpAddress:10.40.40.133
```

SNMP: Support for loopback modes in DS1 MIB

With MAX TNT TAOS 8.0.0, the MAX TNT supports Get and Set requests on the dsx1LoopBackConfig object of the DS1 MIB (RFC 1406).The dsx1NoLoop and dsx1LineLoop loopback modes are supported on both T1 and E1 lines. The dsx1PayloadLoop loopback mode is not supported.

SNMP: Fatal log table

With MAX TNT TAOS 8.0.0, an SNMP user can view the Fatal Log table, and the MAX TNT can inform the user of a reason for a restart when one is available. The following SNMP elements have been introduced to support this feature:

- The fatalLogTable object
- The slotCardResetTrap trap

The sysLastRestartReasonTrap now include fatalLogIndex and sysAbsoluteCurrentTime in its definition:

```
sysLastRestartReasonTrap TRAP-TYPE
    ENTERPRISE    ascend
    VARIABLES     { sysLastRestartReason, fatalLogIndex,
                    sysAbsoluteCurrentTime }
    DESCRIPTION   "This trap is sent to all managers having the
                    alarm condition enabled if the
                    sysLastRestartReason is not unknown
                    (value of 0)."
```

::= 26

A new slotCardResetTrap trap has been defined to inform the MIB manager that a slot card has been reset:

```
slotCardResetTrap TRAP-TYPE
    ENTERPRISE    ascend
    VARIABLES     { fatalLogIndex, fatalLogReason,
                    sysAbsoluteCurrentTime, slotIndex }
    DESCRIPTION   "This trap is sent to all managers having the
                    alarm condition enabled"
```

::= 36

A new table is defined under systemStatusGroup:

```
fatalLogTable OBJECT-TYPE
    SYNTAX        SEQUENCE OF FatalLogTableEntry
    ACCESS        not-accessible
    STATUS        mandatory
```

```
DESCRIPTION "A list of the most recent Fatal Log entries of the
            system as well as the included slot cards"
::= { systemStatusGroup 16 }
```

SNMP: Idle Time variable in the active session table

In previous releases, you could display the idle time for an active session by using the `userstat -o %t` command, but the information was unavailable to SNMP management stations.

With MAX TNT TAOS 8.0.0, the `ssnActiveIdleTime` attribute has been added to the `sessionActiveTable` with Object ID `sessionActiveEntry.8`. The attribute shows the time the session has been idle in 0.01-second increments). Following is the attribute's definition:

```
ssnActiveIdleTime    OBJECT-TYPE
    SYNTAX            TimeTicks
    ACCESS            read-only
    STATUS            mandatory
    DESCRIPTION       "The time, current session has been idle.
                      For non-TNT and non-Max platforms 0 is always
                      reported."
    ::= { sessionActiveEntry 8 }
```

SNMP: Variables added from event.mib to call.mib

With MAX TNT TAOS 8.0.0, a number of new entries have been added to the `callStatus` and `callActive` tables in the Ascend Call MIB, making more information about the call available to the SNMP user.

The `callStatus` Table in the Ascend Call MIB now includes the following new fields:

Field name	Reports
<code>callStatusCalledPartyID</code>	Called party number (if available).
<code>callStatusCallingPartyID</code>	Calling party number (if available). For outgoing calls, this field is set to null.
<code>callStatusMultiLinkID</code>	MP+ bundle ID for MP+ calls. For a non-MP+ call, this field is set to 0 (zero).

The `callActiveTable` in the Ascend Call MIB now includes the following new fields:

Field name	Reports
<code>callActiveCalledParyID</code>	Called party number (if available).
<code>callActiveCallingPartyID</code>	Calling party number (if available). For outgoing calls, this field is set to null.
<code>callActiveMultiLinkID</code>	MP+ bundle ID for MP+ calls. For a non-MP+ call, this field is set to 0 (zero).

SNMP: New MIB variables for summarizing B channel states

The Advanced MIB now includes `wanLineChannelUsageTable(29)`, immediately following `advancedAgent(4)`. The new table contains read-only integer variables that reflect the total count of B channels in any particular state for any given line usage. For example, you can use the table to retrieve the sum of all signaling channels (the number of connected calls) on all trunk lines, or to retrieve the sum of nailed, idle, or ringing channels, or the sum of connected DTPT channels on network (NT) lines.

The new table is indexed by the line usage and B channel state, as defined by the `wanLineUsage` and `wanLineChannelState` enumerations in `advanced.mib`. The MIB currently identifies nine possible line usages and 24 B channel states, yielding a total of 216 new variables that represent the sum of all B channels in a given state for a given line usage.

SNMP: Modified method for adding SNMP object IDs

Previously, algorithms used to assign Object IDs to new MIB members could result in dictionary conflicts across TAOS platforms and software versions. New methodologies make such conflicts much less likely. To ensure that SNMP managers begin using the newer dictionaries that will be maintained across future upgrades, you must compile the new MIB files distributed with MAX TNT TAOS 8.0.0.

SNMP: Details on terminating access resources

SNMP reporting about terminating resources such as modems, HDLC channels, and MultiDSP devices is supported in the Access Resources MIB. In addition, a new trap is supported for notification that a modem is entering a suspect state.

Access Resources MIB

The Ascend Enterprise MIB has been modified to enable the MAX TNT to report utilization and availability details about terminating access resources such as modems, HDLC channels, and MultiDSP devices.

For cards that support `resourceUsageTable` and `resourceTable` in `resource.mib`, the system can report utilization details such as the number of active, available, disabled, suspect, or inoperable devices. This information can be useful for capacity planning and resource management. The system also reports the percentage of available modems, HDLC channels, or DSPs within a device or device group, to enable immediate detection of modem, HDLC, or DSP failure.

The following host cards support this feature:

- Analog Modem (AM36)
- Series56, Series56 II, and Series56 III Digital Modem
- MutiDSP
- Hybrid Access (HDLC2 and HDLC2-EC)

The following object has been added to the Ascend Enterprise MIB (`ascend.mib`):

```
resourcesGroup OBJECT IDENTIFIER ::= { ascend 27 }
```


New trap for modems entering suspect state

The MAX TNT supports a new trap to notify the SNMP management utility when a modem is entering a suspect state. Following is the trap definition in `ascend.trp`:

```
suspectAccessResource TRAP-TYPE
    ENTERPRISE      ascend
    VARIABLES       { resourceSlotIndex, resourcePortIndex,
                     resourceUsedCounts, resourceBadCounts,
                     resourceLast32 }
    DESCRIPTION     "The access resource suspected trap is sent to
                     all the managers in the alarm group when a
                     terminating resource such as modem become suspect.
                     The suspected resource(s) are not assigned to
                     terminate calls until the resource in the available
                     pool exhausted."

::= 34
```

The MAX TNT Trap profile contains the following parameter, shown with its default value:

```
[in TRAP/""]
suspect-access-resource-enabled = no
```

Following is a reference entry for the new parameter.

Suspect-Access-Resource-Enabled

Description: Specifies that whenever a terminating modem has received four successive calls for which it cannot establish a connection, the MAX TNT sends a trap to all SNMP managers in the alarm group.

Once the managing MAX TNT sends the trap, the suspect modem is not assigned to terminate calls until all available resources are exhausted. For example, if a modem drops five calls, the system generates the trap and places the offending modem at the end of the list of available terminating resources.

Usage: Specify one of the following values:

- `yes` directs the MAX TNT to send the `suspectAccessResource` trap when a terminating modem card has received four or more calls for which it could not establish a connection.
- `no` instructs the MAX TNT not to send the `suspectAccessResource` trap.

Example: `set suspect-access-resource-enabled = yes`

Dependencies: The Suspect-Access-Resource-Enabled parameter has an effect only on MAX TNT units with one or more of the following slot cards installed:

- Analog Modem
- Series56, Series56 II, and Series56 III Digital Modem
- MultiDSP

Location: Trap

See Also: Alarm-Enabled, Community-Name, Host-Address, Host-Name, Port-Enabled, Security-Mode

New debug-level command

With debug-level permission, you can access the information from the MAX TNT shelf controller by using the `resrcmgr` command. The command supports the following syntax:

```
admin> ? resrcmgr
usage: resrcmgr -i|u|?
        -i list resource (i)tem information
        -u list resource (u)sage information
        -? display this summary
```

SNMP: Limited support for RFC 2574 user-based security model

With MAX TNT TAOS 8.0.0, MAX TNT units with the network-management license enabled support security enhancements based on the SNMPv3 user-based security model (USM), which is compliant with RFC 2574. The following commands verify that the network management license has been enabled on the system:

```
admin> get base network-management
[in BASE:network-management-enabled]
network-management-enabled = yes
```

With the network-management license, a new SNMPv3-USM-User profile type is available. MAX TNT units support up to 100 configured SNMPv3-USM-User profiles. Configuring the profile enables the USM security features for the specified user.

Note: In this release, encryption is not supported. The Priv-Protocol parameter is set to `no-priv` and its value cannot be modified.

Overview of the SNMPv3 USM settings

Following are the relevant parameters, shown with their default settings:

```
[in SNMPv3-USM-USER/" "]
name* = ""
password = ""
active-enabled = no
read-write-access = no
auth-protocol = md5-auth
priv-protocol = no-priv
```

Parameter	Specifies
Name	Username. Messages sent to or from the SNMP engine on behalf of this name will use the security parameters specified in this profile. The value can contain up to 23 characters and can include special characters by using the <code>\xNN</code> format with the ASCII code for the character. For example, the value <code>test\x20\x21</code> represents the string "test !".
Password	A password, up to 20 characters in length, which maps to a 16 or 20 octet key, in compliance with RFC 2574. Passwords are case sensitive and can include special characters by using the <code>\xNN</code> format with the ASCII code for the character. For example, the value <code>test\x20\x21</code> represents the string "test !" This setting is required if Auth-Protocol is set to a value other than <code>no-auth</code> .

Parameter	Specifies
Active-Enabled	Enable/disable SNMPv3 user-based security model (USM) features for this user. The default value is <code>no</code> .
Read-Write-Access	Enable/disable read-write access to the unit's MIBs for this user. With the default <code>no</code> value, the user has read access only, which enables viewing but not modification of the MIBs.
Auth-Protocol	Enable/disable authentication of messages sent on behalf of this user to or from the SNMP engine, and if enabled, the type of authentication protocol to be used. If this parameter is set to a value other than <code>no-auth</code> , the Password parameter must specify the password to be used. Following are the valid values: <code>no-auth</code> disables authentication for this user. <code>md5-auth</code> (the default value) enables authentication and specifies that the MD5 protocol must be used. <code>sha-auth</code> enables authentication and specifies that SHA protocol must be used.
Priv-Protocol	Enable/disable encryption of messages sent on behalf of this user to or from the SNMP engine, and if enabled, the type of privacy protocol to be used. <i>Not currently supported.</i>

Example of configuring SNMPv3 USM for a user

To configure the USM features for a user, you must specify a name for the profile and set the Active-Enabled parameter to `yes`. You must also specify a password if the Auth-Protocol parameter is set to anything but `no-auth`.

The following commands specify that MD5 authentication must be used for messages sent on behalf of a user named `testv3` to or from the SNMP engine. The user is assigned read-write access to the unit's MIBs.

```
admin> new snmpv3-usm-user testv3
SNMPv3-USM-USER/testv3 read

admin> set password = lma\x2lw
admin> set active-enabled = yes
admin> set read-write-access = yes

admin> write
SNMPv3-USM-USER/testv3 written

admin> list
[in SNMPv3-USM-USER/testv3]
name* = testv3
password = *****
active-enabled = yes
read-write-access = yes
auth-protocol = md5-auth
priv-protocol = no-priv
```

Parameter reference entries

Active-Enabled

Description: Activates a SNMPv3 USM user profile and makes it available for use.

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

Example: `set active-enabled = yes`

Location: SNMPv3-USM-User profile

See Also: Auth-Protocol, Name, Password, Priv-Protocol, Read-Write-Access

Auth-Protocol

Description: Specifies whether or not the MAX TNT unit can authenticate messages sent to and from the SNMP engine, on behalf of the SNMPv3 USM user. Also, specifies the type of authentication protocol the unit uses.

Usage: Specify one of the following settings:

- No-Auth—No authentication.
- MD5-Auth (the default)—The MAX TNT unit uses the MD5 protocol to authenticate incoming and outgoing messages.
- SHA-Auth—The MAX TNT unit uses the SHA protocol to authenticate incoming and outgoing messages.

Example: `set auth-protocol = md5-auth`

Location: SNMPv3-USM-User profile

See Also: Active-Enabled, Name, Password, Priv-Protocol, Read-Write-Access

Name

Description: Specifies the user for whom the MAX TNT unit exchanges an SNMPv3 USM message.

Usage: Specify a name that contains up to 23 characters. The name can include special characters by using the `\xNN` format with the ASCII code for the character. For example, the value `test\x20\x21` represents the string “test !”.

Example: `set name = testv3`

Location: SNMPv3-USM-User profile

See Also: Active-Enabled , Auth-Protocol, Password, Priv-Protocol, Read-Write-Access

Password

Description: Specifies the user’s password, which maps to a 16 or 20 octet key, in compliance with RFC 2574. Passwords are case sensitive.

Usage: Specify up to 20 characters. The password can include special characters by using the `\xNN` format with the ASCII code for the character. For example, the value `test\x20\x21` represents the string “test !”.

Example: `set password = 1rma\x21w`

Dependencies: In the SNMPv3-USM-User profile, you must specify a password if the Auth-Protocol parameter is set to a value other than `no-auth`.

Location: SNMPv3-USM-User profile

See Also: Active-Enabled, Auth-Protocol, Name, Priv-Protocol, Read-Write-Access

Priv-Protocol

Description: Specifies whether or not messages that are sent to or from the SNMP engine can be protected by encryption and the type of privacy protocol to be used.

Usage: The default is `no-priv`. In this release, you cannot change the default setting.

Dependencies: The MAX TNT unit's SNMPv3 engine does not support encryption/decryption.

Location: SNMPv3-USM-User profile

See Also: Active-Enabled, Auth-Protocol, Name, Password, Read-Write-Access

Read-Write-Access

Description: Specifies whether or not the MAX TNT unit grants the SNMPv3 USM user read and write access to the unit's management information base (MIB) settings.

Usage: Specify Yes or No. No is the default. With the default No value, the user has read access only, which enables viewing but not modification of the MIBs.

Example: `set read-write-access = no`

Location: SNMPv3-USM-User profile

See Also: Active-Enabled, Auth-Protocol, Name, Password, Priv-Protocol

IP routing

OSPF nonbroadcast multiaccess (NBMA) support

An OSPF nonbroadcast multiaccess (NBMA) network is any network that has multiple points of access (more than two routers) and does not support broadcast capability. Frame Relay and X.25 are typically NBMA networks.

OSPF routers operate on an NBMA network much as they do on a broadcast network, by using the Hello protocol to form adjacencies and identify the designated router (DR). However, because the routers cannot discover their neighboring routers dynamically by means of broadcasts, you must specify some additional parameters.

With MAX TNT TAOS 8.0.0, the MAX TNT can form adjacencies with other OSPF routers on an NBMA network. Adjacencies enable the unit to route OSPF over Frame Relay networks, and to interoperate with the switches that do not support the serial (point-to-point) model over Frame Relay.

Note: The Non-Multicast parameter in the OSPF-Options subprofiles for IP interfaces causes the translation of the multicast traffic to directed traffic. This parameter is typically used with a serial link, such as a point-to-point connection over Frame Relay, and is not intended for use with NBMA. Non-Multicast must not be enabled for NBMA configurations.

Overview of OSPF NBMA settings

Following are the relevant parameters, shown with default settings:

```
[in IP-INTERFACE/{ { any-shelf any-slot 0 } 0 }:ospf]
network-type = Point-to-Point
poll-interval = 0

[in CONNECTION/"":ip-options:ospf-options]
network-type = Point-to-Point
poll-interval = 0

[in OSPF-NBMA-NEIGHBOR/" " (new)]
name* = ""
host-name = ""
ip-address = 0.0.0.0
dr-capable = no
```


Parameter	Specifies
Network-Type	The kind of network to which the interface connects: Broadcast, NonBroadcast (multiaccess), or Point-to-Point. Broadcast is any broadcast-capable network, such as Ethernet. NonBroadcast (multiaccess) is used for networks that have more than two OSPF routers and no broadcast capability, such as Frame Relay or X.25. Point-to-Point (the default), is used for interfaces connected to one other node on the other far end.
Poll-Interval	Interval in seconds at which to send Hello packets to a neighboring router that has become inactive. The default 0 (zero) means that no Hello packets are sent to a neighboring router from which no Hello packets have been received for the number of seconds specified in the Dead-Interval setting. If you specify a nonzero value, use a larger value than the regular Hello-Interval default of 10 seconds (for example, 120 seconds).
Name	The name of the OSPF-NBMA-Neighbor profile.
Host-Name	Station name of a local Connection profile that defines the connection to the neighboring router.
IP-Address	IP address of the neighboring router.
DR-Capable	Indication of whether the neighboring router can be the designated router (DR). Values are <i>yes</i> and <i>no</i> (the default).

Example of an OSPF NBMA configuration

On an NBMA network, a router that is eligible to become the DR is configured with a list of all other OSPF routers connected to the network. At startup, these routers send Hello packets to each other to discover the DR. The DR then begins sending Hello packets to the entire list of routers on the network. When an NBMA interface becomes active on the MAX TNT, it sends Hello packets only to neighboring routers that are eligible to become the DR, until it is notified about which router is the DR.

Figure 12 shows an OSPF NBMA network using Frame Relay. For the purposes of this example, assume that the system named FR-Router is eligible to become the DR, and that the MAX-Router unit is not DR-capable.

Figure 12. OSPF nonbroadcast multiaccess (NBMA) network



Example of configuring a DR-capable neighboring router

The following commands define a sample Frame-Relay profile for the interface to FR-Router in Figure 12:

```
admin> new frame-relay fr-dce
FRAME-RELAY/fr-dce read
admin> set active = yes
admin> set link-type = dce
admin> set nailed-up-group = 36
admin> set link-mgmt = ccitt
admin> write
FRAME-RELAY/fr-dce written
```

The next commands define a Connection profile to FR-Router:

```
admin> new conn FR-Router
[ in CONNECTION/FR-Router (new) ]
admin> set active = yes
admin> set encapsulation-protocol = frame-relay
admin> set ip-options remote-address = 10.1.1.1/24
admin> set ip-options ospf active = yes
admin> set ip-options ospf area = 0.0.0.1
admin> set ip-options ospf network-type = NonBroadcast
admin> set ip-options ospf poll-interval = 60
admin> set telco-options call-type = ft1
admin> set fr-options frame-relay-profile = fr-dce
admin> set fr-options dlci = 100
admin> write
CONNECTION/FR-Router written
```

The next commands enable the MAX TNT to form an adjacency with FR-Router:

```
admin> new ospf-nbma-neighbor fr-router
[in OSPF-NBMA-NEIGHBOR/fr-router (new)]

admin> set host-name = FR-Router

admin> set ip-address = 10.1.1.1/24

admin> set dr-capable = yes

admin> write
OSPF-NBMA-NEIGHBOR/fr-router written
```

Example of configuring a non-DR-capable neighbor

The following commands define a Frame-Relay profile for link operations on the interface to the system named MAX-Router in Figure 12:

```
admin> new frame-relay fr-dte
FRAME-RELAY/fr-dte read

admin> set active = yes

admin> set link-type = dte

admin> set nailed-up-group = 11

admin> set link-mgmt = ccitt

admin> write
FRAME-RELAY/fr-dte written
```

The next commands define a Connection profile to MAX-Router:

```
admin> new conn MAX-Router
[in CONNECTION/MAX-Router (new)]

admin> set active = yes

admin> set encapsulation-protocol = frame-relay

admin> set ip-options remote-address = 20.2.2.2/28

admin> set ip-options ospf active = yes

admin> set ip-options ospf area = 0.0.0.1

admin> set ip-options ospf network-type = NonBroadcast

admin> set ip-options ospf poll-interval = 60

admin> set telco-options call-type = ft1

admin> set fr-options frame-relay-profile = fr-dte

admin> set fr-options dlci = 200

admin> write
CONNECTION/MAX-Router written
```

The next commands enable the MAX TNT to form an adjacency with FR-Router:

```
admin> new ospf-nbma-neighbor max-router
[in OSPF-NBMA-NEIGHBOR/max-router (new)]

admin> set host-name = MAX-Router

admin> set ip-address = 20.2.2.2/28

admin> write
OSPF-NBMA-NEIGHBOR/max-router written
```

OSPF area border router (ABR) support

In previous releases, the MAX TNT acted as an OSPF internal router with limited border router capability, so its LAN and WAN interfaces all had to be in the same area and area-type. With MAX TNT TAOS 8.0.0, MAX TNT units support OSPF ABR for all area types, including not-so-stubby areas (NSSAs). For details about ABR functionality, see RFC 2328, *OSPF Version 2*. For information about NSSA ABR operations, see RFC 1587, *The OSPF NSSA Option*.

NSSAs are like stub areas in that they do not receive or originate Type-5 link state advertisements (LSAs). However, NSSAs employ Type-7 LSAs for carrying ASE route information within the area. Type 7 LSAs use a propagate (P) bit to flag the NSSA border router to translate the Type-7 LSA into a type-5 LSA, which can then be propagated into other areas. The following new command lists the router IDs of NSSA border routers (which are performing the Type-7 to Type-5 LSA translation):

```
admin> ospf translators

Area ID      Router ID
0.0.0.1      10.105.0.13
0.0.0.2      12.1.1.1
```

OSPF MD5 16-byte authentication key

In a previous release, MAX TNT units added support for the MD5 cryptographic authentication method for OSPF protocol exchanges. You set the secret key to be used for MD5 authentication by using the Auth-Key parameter, which is also used for null or simple authentication. However, the Auth-Key parameter accepts a maximum of only 8 characters.

With MAX TNT TAOS 8.0.0, new parameters are introduced to support a secret authentication key of up to 16 characters. This change is in compliance with RFC 2328, *OSPF Version 2*. Following are the new parameters, shown with default settings:

```
[in IP-INTERFACE/{ { any-shelf any-slot 0 } 0 }:ospf]
md5-auth-key = *****

[in CONNECTION/" ":ip-options:ospf-options]
md5-auth-key = *****

[in OSPF-VIRTUAL-LINK/0/0/0/0]
md5-authen-key = *****
```

Parameter	Specifies
MD5-Auth-Key	Secret key to be used for the MD5 cryptographic authentication method, up to 16 characters. The default value is <code>ascend0</code> . When Authen-Type is set to <code>md5</code> , you must supply a key in the new field because the Auth-Key setting used previously no longer applies.
MD5-Authen-Key	Secret key to be used for the MD5 cryptographic authentication method, up to 16 characters. The default value is <code>ascend0</code> . When Authen-Type is set to <code>md5</code> , you must supply a key in the new field because the Auth-Key setting used previously no longer applies.

For example, the following commands enable OSPF routing, specify MD5 authentication and supply a key on a LAN interface:

```

admin> read ip-interface { { 1 12 1 } 0 }
IP-INTERFACE/{ { shelf-1 slot-12 1 } 0 } read

admin> set ospf active = yes

admin> set ospf authen-type = md5

admin> set ospf md5-auth-key = 12!secret*34key

admin> write
IP-INTERFACE/{ { shelf-1 slot-12 1 } 0 } written

```

RFC 1850 OSPF traps

With MAX TNT TAOS 8.0.0, MAX TNT units support OSPF traps as defined in RFC 1850, *OSPF Version 2 Management Information Base*. For an OSPF trap to be generated when the trap condition occurs, OSPF traps must be enabled, either in the Trap profile or by setting the corresponding bit in the new MIB object, `ospfSetTrap`, defined in RFC 1850. In addition, the individual trap that represents the trap condition must be enabled.

Overview of trap definitions

Following are the relevant parameters (shown with default values) in a Trap profile:

```

[in TRAP/""]
ospf-enabled = no
ospf-if-config-error-enabled = no
ospf-if-auth-failure-enabled = no
ospf-if-state-change-enabled = no
ospf-if-rx-bad-packet = no
ospf-tx-retransmit-enabled = no
ospf-nbr-state-change-enabled = no
ospf-virt-if-config-error-enabled = no
ospf-virt-if-auth-failure-enabled = no
ospf-virt-if-state-change-enabled = no
ospf-virt-if-rx-bad-packet = no
ospf-virt-if-tx-retransmit-enabled = no
ospf-virt-nbr-state-change-enabled = no
ospf-originateLsa-enabled = no
ospf-maxAgeLsa-enabled = no
ospf-lsdb-overflow-enabled = no
ospf-approaching-overflow-enabled = no

```

Parameter	Specifies
OSPF-enabled	Enable/disable generation of OSPF traps. When set to <code>no</code> (the default), no OSPF traps are generated regardless of individual OSPF trap settings in the profile. When set to <code>yes</code> , trap generation depends on whether the specific OSPF trap is enabled.
OSPF-if-config-error-enabled	Enable/disable trap generation if a packet has been received on a nonvirtual interface from a router whose configuration conflicts with this router's configuration. The system generates this trap when it detects configuration error types from 1 to 9, as defined in RFC 1850. Generation of the trap typically indicates a failure to form an adjacency, although this is not always the case. Traps for error type 10 (optionsMismatch) are not currently supported. (OSPF Trap 4)

Parameter	Specifies
OSPF-if-auth-failure-enabled	Enable/disable trap generation if a packet has been received on a nonvirtual interface from a router whose authentication key or authentication type conflicts with this router's authentication key or authentication type. (OSPF Trap 6)
OSPF-if-state-change-enabled	Enable/disable trap generation if the state of a nonvirtual OSPF interface has changed. This trap is generated when the interface state regresses (for example, goes from Dr to Down) or progresses to a terminal state (Point-to-Point, DR Other, Dr, or Backup). (OSPF Trap 16)
OSPF-if-rx-bad-packet	Enable/disable trap generation if an OSPF packet has been received on a nonvirtual interface that cannot be parsed. (OSPF Trap 8)
OSPF-tx-retransmit-enabled	Enable/disable trap generation if an OSPF packet has been retransmitted on a nonvirtual interface. All packets that are retransmitted are associated with a link-state database (LSDB) entry. The LS type, LS ID, and Router ID are used to identify the LSDB entry. (OSPF Trap 10)
OSPF-nbr-state-change-enabled	Enable/disable trap generation if the state of a nonvirtual OSPF neighbor has changed. This trap is generated when the neighbor state regresses (for example, changes from Attempt or Full to 1-Way or Down) or progresses to a terminal state (for example, 2-Way or Full). When an neighbor transitions from or to Full on nonbroadcast multiaccess (NBMA) and broadcast networks, the trap is generated by the designated router. A designated router transitioning to Down is noted by OSPFIfStateChange. (OSPF Trap 2)
OSPF-virt-if-config-error-enabled	Enable/disable trap generation if a packet has been received on a virtual interface from a router whose configuration parameters conflict with this router's configuration parameters. The system generates this trap when it detects configuration error types from 1 to 9, as defined in RFC 1850. Generation of the trap typically indicates a failure to form an adjacency, although this is not always the case. Traps for error type 10 (optionsMismatch) are not currently supported. (OSPF Trap 5)
OSPF-virt-if-auth-failure-enabled	Enable/disable trap generation if a packet has been received on a virtual interface from a router whose authentication key or authentication type conflicts with this router's authentication key or authentication type. (OSPF Trap 7)
OSPF-virt-if-state-change-enabled	Enable/disable trap generation if the state of an OSPF virtual interface has changed. (OSPF Trap 1)
OSPF-virt-if-rx-bad-packet	Enable/disable trap generation if an OSPF packet has been received on a virtual interface that cannot be parsed. (OSPF Trap 9)
OSPF-virt-if-tx-retransmit-enabled	Enable/disable trap generation if an OSPF packet has been retransmitted on a virtual interface. All packets that are retransmitted are associated with an LSDB entry. The LS type, LS ID, and Router ID are used to identify the LSDB entry. (OSPF Trap 11)

Parameter	Specifies
OSPF-if-auth-failure-enabled	Enable/disable trap generation if a packet has been received on a nonvirtual interface from a router whose authentication key or authentication type conflicts with this router's authentication key or authentication type. (OSPF Trap 6)
OSPF-if-state-change-enabled	Enable/disable trap generation if the state of a nonvirtual OSPF interface has changed. This trap is generated when the interface state regresses (for example, goes from Dr to Down) or progresses to a terminal state (Point-to-Point, DR Other, Dr, or Backup). (OSPF Trap 16)
OSPF-if-rx-bad-packet	Enable/disable trap generation if an OSPF packet has been received on a nonvirtual interface that cannot be parsed. (OSPF Trap 8)
OSPF-tx-retransmit-enabled	Enable/disable trap generation if an OSPF packet has been retransmitted on a nonvirtual interface. All packets that are retransmitted are associated with a link-state database (LSDB) entry. The LS type, LS ID, and Router ID are used to identify the LSDB entry. (OSPF Trap 10)
OSPF-nbr-state-change-enabled	Enable/disable trap generation if the state of a nonvirtual OSPF neighbor has changed. This trap is generated when the neighbor state regresses (for example, changes from Attempt or Full to 1-Way or Down) or progresses to a terminal state (for example, 2-Way or Full). When an neighbor transitions from or to Full on nonbroadcast multiaccess (NBMA) and broadcast networks, the trap is generated by the designated router. A designated router transitioning to Down is noted by OSPFIfStateChange. (OSPF Trap 2)
OSPF-virt-if-config-error-enabled	Enable/disable trap generation if a packet has been received on a virtual interface from a router whose configuration parameters conflict with this router's configuration parameters. The system generates this trap when it detects configuration error types from 1 to 9, as defined in RFC 1850. Generation of the trap typically indicates a failure to form an adjacency, although this is not always the case. Traps for error type 10 (optionsMismatch) are not currently supported. (OSPF Trap 5)
OSPF-virt-if-auth-failure-enabled	Enable/disable trap generation if a packet has been received on a virtual interface from a router whose authentication key or authentication type conflicts with this router's authentication key or authentication type. (OSPF Trap 7)
OSPF-virt-if-state-change-enabled	Enable/disable trap generation if the state of an OSPF virtual interface has changed. (OSPF Trap 1)
OSPF-virt-if-rx-bad-packet	Enable/disable trap generation if an OSPF packet has been received on a virtual interface that cannot be parsed. (OSPF Trap 9)
OSPF-virt-if-tx-retransmit-enabled	Enable/disable trap generation if an OSPF packet has been retransmitted on a virtual interface. All packets that are retransmitted are associated with an LSDB entry. The LS type, LS ID, and Router ID are used to identify the LSDB entry. (OSPF Trap 11)

Parameter	Specifies
OSPF-virt-nbr-state-change-enabled	Enable/disable trap generation if the state of an OSPF virtual neighbor has changed. (OSPF Trap 3)
OSPF-originateLsa-enabled	Enable/disable trap generation if a new LSA has been originated by this router due to a topology change. (OSPF Trap 12)
OSPF-maxAgeLsa-enabled	Enable/disable trap generation if an LSA in the router's link-state database has aged to MaxAge. (OSPF Trap 13)
OSPF-lsdb-overflow-enabled	Enable/disable trap generation if the number of LSAs in the router's link-state database has exceeded OSPFExtLsdbLimit. (OSPF Trap 14)
OSPF-approaching-overflow-enabled	Enable/disable trap generation if the number of LSAs in the router's link-state database has exceeded 90 percent of OSPFExtLsdbLimit. (OSPF Trap 15)

Example of setting traps in the Trap profile

The following commands cause the system to generate traps when the router receives a packet from an OSPF router in which a configuration mismatch (such as an invalid OSPF version number or an address conflict) or an authentication failure occurs:

```
admin> read trap monitor-ospf
TRAP/monitor-ospf read
admin> set ospf-enabled = yes
admin> set ospf-if-config-error-enabled = yes
admin> set ospf-if-auth-failure-enabled = yes
admin> write
TRAP/monitor-ospf written
```

SNMP support for OSPF traps

In addition to the Trap profile changes, a new MIB (`rfc1850.mib`) is now distributed as part of this release. Management stations and browsers used to manage OSPF should now load `rfc1850.mib` instead of the old `rfc1253.mib`. A new MIB object, `ospfSetTrap` is defined according to RFC 1850 for enabling trap events:

```
.iso.org.dod.internet.mgmt.mib-2. ospf.ospfTrap.ospfTrapControl.ospfSetTrap
```

This object defaults initially to the octet string `{ '\0x0', '0x0', '0x0', '0x0' }` (or the hex value 0x0), which disables all trap events. The value of this object is stored in NVRAM.

OSPF reconfiguration restart no longer required

In previous releases, changes to the OSPF configuration required a system reset to take effect. With MAX TNT TAOS 8.0.0, a system reset is no longer required for this purpose. When you write a profile with configuration changes that affect OSPF, OSPF drops its adjacencies and reinitializes with the new configuration values.

New parameter to disable OSPF

To enable you to globally disable the OSPF protocol, the system supports the following new parameter, shown with its default value:

```
[in IP-GLOBAL:ospf-global]
enable = yes
```

Parameter	Specifies
Enable	Starts or stops the OSPF protocol. A change to the setting takes effect immediately upon writing the profile.

In previous software versions, the only way to globally deactivate OSPF was to disable it manually on each OSPF interface. This parameter provides a global mechanism for disabling the protocol. It can also be used to prevent OSPF from reinitializing several times if you are modifying many OSPF-related profiles. In that case, set the parameter to `no`, write the OSPF changes, and then set the parameter to `yes` again.

Changes to existing administrative commands

The `ospfd` and `ospf` commands have been modified to be usable when OSPF is disabled. In addition, a new `log` option has been added to the `ospfd` command to write its trace messages as debug-level log messages. This option allows the system to collect traces while OSPF initializes.

Per-connection Microsoft WINS assignment

In previous releases, the MAX TNT allowed system-wide configuration of a primary and secondary NetBIOS Windows Internet Name Service (WINS) server, to support WINS name resolution for machines connected to a NetBIOS network.

With MAX TNT TAOS 8.0.0, you can specify a primary and secondary WINS server on a per-connection basis, either in local Connection profiles or in RADIUS.

Note: The PC dialing in must have the Dynamic Host Configuration Protocol (DHCP) for WINS enabled in its Network settings for this feature to work.

Settings in a Connection profile

Following are the local parameters (shown with default settings) for configuring client WINS servers:

```
[in CONNECTION/" ":ip-options]
client-wins-primary-addr = 0.0.0.0
client-wins-secondary-addr = 0.0.0.0
client-wins-addr-assign = yes
```

Parameter	Specifies
Client-WINS-Primary-Addr	Address of a client WINS server for the connection.

Client-WINS-Secondary-Addr	Address of a secondary client WINS server for the connection.
Client-WINS-Addr-Assign	Enable/disable client WINS for the connection. If set to <code>true</code> (the default), the system presents client WINS server addresses while negotiating the connection.

For more details about these parameters, see “Parameter reference entries” on page 144. For information about specifying NetBIOS servers in the IP-Global profile, see the *MAX TNT Network Configuration Guide*.

Settings in a RADIUS profile

The following attribute-value pairs configure client WINS servers in RADIUS profiles:

RADIUS attribute	Value
Ascend-Client-Primary-WINS (78)	Address of a client WINS server for the connection.
Ascend-Client-Secondary-WINS (79)	Address of a secondary client WINS server for the connection.
Ascend-Client-Assign-WINS (80)	Enable/disable the use of client WINS servers for the connection. If set to <code>WINS-Assign-Yes (1)</code> , the system presents client WINS server addresses while negotiating the connection.

For more details about these attributes, see “RADIUS attribute reference entries” on page 145.

To use these attributes, the RADIUS server must support vendor-specific attributes (VSAs) and the MAX TNT must be configured in VSA compatibility mode. Following are the relevant settings:

```
[in EXTERNAL-AUTH]
auth-type = radius

[in EXTERNAL-AUTH:rad-auth-client]
auth-radius-compat = vendor-specific
```

For details about these settings, see the *MAX TNT Reference Guide*.

Examples of configuring client WINS servers

The following commands identify two WINS servers for this connection. The secondary server is accessed only if the primary one is inaccessible.

```
admin> read connection pc-1
CONNECTION/pc-1 read

admin> set ip-options client-wins-primary-addr = 10.2.3.4
admin> set ip-options client-wins-secondary-addr = 10.2.3.56
admin> set ip-options client-wins-addr-assign = yes
admin> write
CONNECTION/pc-1 written
```

Following are comparable settings in a RADIUS profile:

```
pc-1 Password = "localpw", Service-Type = Framed
    Framed-Protocol = PPP,
    Framed-IP-Address = 1.1.1.1,
    Framed-IP-Netmask = 255.255.255.0,
    Ascend-Client-Primary-WINS = 10.2.3.4,
    Ascend-Client-Secondary-WINS = 10.2.3.56,
    Ascend-Client-Assign-WINS = WINS-Assign-Yes
```

Parameter reference entries

Client-WINS-Addr-Assign

Description: Specifies whether the MAX TNT presents client WINS server addresses to the dial-in client while negotiating the session.

Usage: Specify *yes* (the default) or *no*. A *no* setting still enables the PC to access WINS name resolution if NetBIOS servers have been configured in the IP-Global profile.

Example: `set client-wins-addr-assign = yes`

Dependencies: The PC dialing in must have DHCP for WINS enabled in its Network settings for the client WINS feature to work. The IP addresses of one or more WINS servers must be specified in the Client-WINS-Primary-Addr and (optionally) Client-WINS-Secondary-Addr parameters in the Connection profile.

Location: Connection > IP-Options

See Also: Client-WINS-Primary-Addr, Client-WINS-Secondary-Addr, NetBIOS-Primary-Ns, NetBIOS-Secondary-Ns

Client-WINS-Primary-Addr

Description: Specifies the IP address of the primary WINS server. The primary server is used for WINS name resolution. The secondary server, if one is specified, is used only if the primary server is unavailable.

Usage: Specify the IP address of a WINS server.

Example: `set client-wins-primary-addr = 10.1.1.1`

Dependencies: The PC dialing in must have DHCP for WINS enabled in its Network settings for the client WINS feature to work. Client-WINS-Addr-Assign must be set to *yes* for the server address to be passed to the dial-in client during session negotiation.

Location: Connection > IP-Options

See Also: Client-WINS-Addr-Assign, Client-WINS-Secondary-Addr, NetBIOS-Primary-Ns, NetBIOS-Secondary-Ns

Client-WINS-Secondary-Addr

Description: Specifies the IP address of the secondary WINS server. The secondary server is used for WINS name resolution only if the primary server is unavailable.

Usage: Specify the IP address of a WINS server.

Example: `set client-wins-secondary-addr = 20.1.1.1`

Dependencies: The PC dialing in must have DHCP for WINS enabled in its Network settings for the client WINS feature to work. Client-WINS-Addr-Assign must be set to *yes* for the server address to be passed to the dial-in client during session negotiation.

Location: Connection > IP-Options

See Also: Client-WINS-Addr-Assign, Client-WINS-Primary-Addr, NetBIOS-Primary-Ns, NetBIOS-Secondary-Ns

RADIUS attribute reference entries

Ascend-Client-Primary-WINS (78)

Description: Specifies the IP address of the primary WINS server. The primary server is used for WINS name resolution. The secondary server, if one is specified, is used only if the primary server is unavailable.

Usage: Specify the IP address of a WINS server.

Example: `Ascend-Client-Primary-WINS = 10.1.1.1`

Dependencies: The PC dialing in must have DHCP for WINS enabled in its Network settings for the client WINS feature to work. Ascend-Client-Assign-WINS must be set to WINS-Assign-Yes (1) for the server address to be passed to the dial-in client during session negotiation. To use this feature, you must set the Auth-Compat-Mode parameter on the MAX TNT to Vendor-Specific.

See Also: Ascend-Client-Secondary-WINS (79), Ascend-Client-Assign-WINS (80)

Ascend-Client-Secondary-WINS (79)

Description: Specifies the IP address of the secondary WINS server. The secondary server is used for WINS name resolution only if the primary server is unavailable.

Usage: Specify the IP address of a WINS server.

Example: `Ascend-Client-Secondary-WINS = 20.1.1.1`

Dependencies: The PC dialing in must have DHCP for WINS enabled in its Network settings for the client WINS feature to work. Ascend-Client-Assign-WINS must be set to WINS-Assign-Yes (1) for the server address to be passed to the dial-in client during session negotiation. To use this feature, you must set the Auth-Compat-Mode parameter on the MAX TNT to Vendor-Specific.

Dependencies: Ascend-Client-Primary-WINS (78), Ascend-Client-Assign-WINS (80)

Ascend-Client-Assign-WINS (80)

Description: Specifies whether the MAX TNT presents client WINS server addresses to the dial-in client while negotiating the session.

Usage: Specify WINS-Assign-Yes (1) or WINS-Assign-No (0). A setting of WINS-Assign-No (0) still enables the PC dialing in to access WINS name resolution if NetBIOS servers have been configured in the IP-Global profile.

Example: `Ascend-Client-Assign-WINS = 1`

Dependencies: The PC dialing in must have DHCP for WINS enabled in its Network settings for the client WINS feature to work. The IP addresses of one or more WINS servers must be specified in attributes 78 and (optionally) 79.

See Also: Ascend-Client-Primary-WINS (78), Ascend-Client-Secondary-WINS (79)

Private routing tables

With MAX TNT TAOS 8.0.0, you can define private routing tables. Only Connection or RADIUS profiles that refer to the private routing table have access to its route definitions.

In previous releases, private routes could be defined in a RADIUS user profile through the use of the Ascend-Private-Route (104) attribute. Now, you can also use that attribute in private-route pseudo-user profiles, which can then be referred to by multiple RADIUS or Connection profiles, or both. These externally defined private tables are cached locally for a configurable interval. The PrtCache command displays statistics about each cached RADIUS private-route profile, and enables you to flush profiles from the cache.

You can also define private routing tables locally, in a new Private-Route-Table profile. These profiles can then be referenced by multiple RADIUS or Connection profiles, or both.

Overview of local profile settings

```
[in PRIVATE-ROUTE-TABLE/" " ]
name* = " "

[in PRIVATE-ROUTE-TABLE/" ":route-description-list[1]]
enabled = no
dest-address = 0.0.0.0/0
netmask = 0.0.0.0
gateway-address = 0.0.0.0
metric = 0

[in CONNECTION/" ":ip-options]
private-route-table = " "
private-route-profile-required = no

[in ANSWER-DEFAULTS:ip-answer]
private-route-profile-required = no

[in IP-GLOBAL]
default-prt-cache-time = 1440
```

Parameter	Specifies
Name	Name of the profile, up to 23 characters. This name is used to associate a RADIUS or Connection profile with the defined private routes.
Enabled	Enable/disable the specific route for use in the private routing table. A table can contain up to 24 routes.
Dest-Address	Destination IP address, which can include a subnet specification. This setting works the same as its counterpart in an IP-Route profile. For details, see the <i>MAX TNT Reference Guide</i> .
Netmask	Netmask of the destination IP address, set automatically when you specify a prefix length as part of the IP address.

Parameter	Specifies
Gateway-Address	IP address of a router used to reach the specified destination. This setting works the same as its counterpart in an IP-Route profile. For details, see the <i>MAX TNT Reference Guide</i> .
Metric	RIP metric for the specified route (a number between 1 to 15; the default is 8). This setting works the same as its counterpart in an IP-Route profile. For details, see the <i>MAX TNT Reference Guide</i> .
Private-Route-Table	Name of a Private-Route-Table profile associated with the connection. The name can be of a local profile or a private-route pseudo-user profile in RADIUS. However, if a local Connection profile does not use authentication, it cannot point to a RADIUS private-route profile.
Private-Route-Profile-Required	Whether access to the private routing table is required for the session. With the default value of No, the system establishes the session even if the associated private routing table is not found. If the parameter is set to yes, the system disconnects the call if the specified private routing table is not found. This setting does not apply if the profile does not refer to a private routing table by name. The Answer-Defaults setting is used for RADIUS user profiles that refer to a private routing table and do not explicitly specify a value for Ascend-Private-Route-Required (55).
Default-Prt-Cache-Time	Number of minutes to cache RADIUS private route profiles that do not include a value for Ascend-Cache-Time (57). The default is 1440 (24 hours). Once the cache timer expires, cached profiles are deleted from system memory. The next time a private route is needed, the system retrieves the profile from RADIUS and stores it in cache again. Keeping a profile in cache increases the performance of route lookups, at the cost of some system memory. If this parameter is set to 0 (zero), the default timer is disabled so that only RADIUS profiles specifying a cache time are cached.

Overview of RADIUS attributes to refer to a private route table

RADIUS user profiles can refer to private-route profiles by specifying the following vendor-specific attributes (VSAs):

RADIUS attribute	Value
Ascend-Private-Route-Table-ID (54)	Name of a RADIUS private-route profile associated with the connection.
Ascend-Private-Route-Required (55)	Whether access to the private routing table is required for the session. With the default value of Required-No (0), the system establishes the session even if the associated private routing table is not found. If the attribute is set to Required-Yes (1), the system disconnects the call if the private routing table is not found. This setting does not apply if the profile does not refer to a private routing table by name. If this attribute is not specified, the Answer-Defaults setting is used.

Overview of RADIUS attributes to define a private route table

In RADIUS, private route tables are defined in a pseudo-user profile. A private-route profile is a pseudo-user profile in which the first two lines have the following format:

```
profile-name Password = "ascend" Service-Type = Outbound
```

The *profile-name* value is any name you assign to the profile. Private-route profile definitions can include the following VSAs:

RADIUS attribute	Value
Ascend-Private-Route (104)	Destination address and next-hop router address for a private route. Each private-route profile specifies one or more private routes with the Ascend-Private-Route (104) attribute, which is described in the <i>MAX TNT RADIUS Guide</i> .
Ascend-Cache-Refresh (56)	Whether the timer for cached routes in this profile is reset each time a new session becomes active that refers to the pseudo-user profile. Refresh-No (0) does not reset the timer. Refresh-Yes (1) resets the cache timer when a session referring to the profile becomes active.
Ascend-Cache-Time (57)	Number of minutes to cache the profile. Once the cache timer expires for a RADIUS profile, the profile is deleted from system memory. The next time it is needed, the system retrieves it from RADIUS and stores it in cache again. Keeping a profile in cache increases the performance of route lookups, at the cost of some system memory. The minimum possible cache time is 0 minutes, which causes the system to retrieve the profile for every route lookup in the table. This setting is usually not desirable. If this attribute is not specified, the IP-Global setting is used.

To use these attributes, the RADIUS server must support vendor-specific attributes (VSAs) and the MAX TNT must be configured in VSA compatibility mode. Following are the relevant settings:

```
[in EXTERNAL-AUTH]
auth-type = radius

[in EXTERNAL-AUTH:rad-auth-client]
auth-radius-compatible = vendor-specific
```

For details about these settings, see the *MAX TNT Reference Guide*.

Examples of configuring a private routing table

You can configure private routing tables locally or in RADIUS. For example, the following commands define a private routing table named *check*:

```
admin> new private-route-table check
PRIVATE-ROUTE-TABLE/check read

admin> list route-description-list 1
[in PRIVATE-ROUTE-TABLE/check:route-description-list[1] (new)]
enabled = no
dest-address = 0.0.0.0/0
netmask = 0.0.0.0
```

```
gateway-address = 0.0.0.0
metric = 0

admin> set enabled = yes

admin> set dest-address = 1.1.1.1/24

admin> set gateway-address = 2.2.2.2

admin> set metric = 2

admin> list
[in PRIVATE-ROUTE-TABLE/check:route-description-list[1]]
enabled = yes
dest-address = 1.1.1.1/24
netmask = 255.255.255.0
gateway-address = 2.2.2.2
metric = 2

admin> list .. 2
[in PRIVATE-ROUTE-TABLE/check:route-description-list[1]]
enabled = no
dest-address = 0.0.0.0/0
netmask = 0.0.0.0
gateway-address = 0.0.0.0
metric = 0

admin> set enabled = yes

admin> set dest-address = 3.3.3.3/28

admin> set gateway-address = 2.2.2.2

admin> set metric = 3

admin> write
PRIVATE-ROUTE-TABLE/check written
```

Following is a comparable RADIUS private-route profile:

```
check Password = "ascend", Service-Type = Outbound
  Ascend-Cache-Time = 3,
  Ascend-Cache-Refresh = Refresh-Yes,
  Ascend-Private-Route = "1.1.1.1/24 2.2.2.2 2",
  Ascend-Private-Route = "3.3.3.3/28 2.2.2.2 3"
```

The following commands configure the default cache time for RADIUS private-route profiles:

```
admin> read ip-global
IP-GLOBAL read

admin> set default-prt-cache-time = 180

admin> write
IP-GLOBAL written
```

Following is a sample RADIUS private-route profile that makes use of the default because a value for Ascend-Cache-Time (57) is not explicitly specified:

```
my-routes Password = "ascend"
  Service-Type = Outbound,
  Ascend-Private-Route = "1.1.1.1/24 2.2.2.2",
  Ascend-Private-Route = "3.3.3.3/28 2.2.2.2"
```

Examples of using private routing tables

The following commands modify a Connection profile so that the session has access to those routes in the private table, and the system disconnects the call if the private table is not found:

```
admin> read connection p50-v2
CONNECTION/p50-v2 read

admin> set ip-options private-route-table = check

admin> set ip-options private-route-profile-required = yes

admin> write
CONNECTION/p50-v2 written
```

Following is a sample RADIUS profile that refers to the same private table with the same requirements. This profile also specifies how the routes are cached for this connection:

```
p50-v2 Password = "my-password"
  Service-Type = Framed,
  Framed-Protocol = PPP,
  Framed-IP-Address = 10.1.1.1,
  Framed-IP-Netmask = 255.0.0.0,
  Ascend-Private-Route-Table-ID = "check",
  Ascend-Private-Route-Required = Required-Yes
```

The following commands configure the system to reject incoming calls when the RADIUS user profile specifies a private routing table that is not found:

```
admin> read answer-defaults
ANSWER-DEFAULTS read

admin> set ip-answer private-route-profile-required = yes

admin> write
ANSWER-DEFAULTS written
```

Following is a sample RADIUS profile that makes use of the default because a value for Ascend-Private-Route-Required (55) is not explicitly specified:

```
p50-v2 Password = "my-password"
  Service-Type = Framed,
  Framed-Protocol = PPP,
  Framed-IP-Address = 10.1.1.1,
  Framed-IP-Netmask = 255.0.0.0,
  Ascend-Private-Route-Table-ID = "check"
```

Parameter reference entries

Default-Prt-Cache-Time

Description: Specifies the default cache time for private-route profiles configured in RADIUS.

Usage: Specify an integer in minutes. The default is 1440 minutes. If you set Default-Prt-Cache-Time to 0 (zero), RADIUS private-route profiles are not cached.

Example: `set default-prt-cache-time = 1200`

Dependencies: The system uses the Default-Prt-Cache-Time value only if no cache time is specified in the RADIUS private-route profile.

Location: IP-Global

Private-Route-Profile-Required

Description: In the Answer-Defaults profile, specifies whether the system drops the call if it cannot locate the Private-Route-Table profile indicated in the RADIUS user profile. In a Connection profile, this parameter specifies whether the system drops the call if it cannot locate the Private-Route-Table profile indicated in the Connection profile.

Usage: Specify *yes* or *no*. The default is *no*.

- *yes* specifies that the system drops the call if it cannot locate the Private-Route-Table profile.
- *no* specifies that the system establishes the link even if it cannot locate the Private-Route-Table profile.

Example: `set private-route-profile-required = yes`

Dependencies: The unit uses the Private-Route-Profile-Required value in the Answer-Defaults profile only if the Ascend-Private-Route-Required attribute is not set in a RADIUS private-route profile.

Location: Answer-Defaults > IP-Answer, Connection > IP-Options

See Also: Private-Route-Table

Private-Route-Table

Description: Specifies the private routing table for the connection.

Usage: Specify the name of the Private-Route-Table profile associated with the connection. You can enter up to 23 characters. The default is *null*.

Example: `set private-route-table = private-rt-1`

Location: Connection > IP-Options

See Also: Private-Route-Profile-Required

RADIUS attribute reference entries

Ascend-Private-Route-Table-ID (54)

Description: Specifies the name of the private-route profile associated with the connection. The table can be defined in a profile on the MAX TNT or in a RADIUS pseudo-user profile.

Usage: Specify a text string.

See Also: Ascend-Private-Route-Required

Ascend-Private-Route-Required (55)

Description: Specifies whether a connection can be established if its associated private-route profile is not found.

Usage: Specify one of the following values:

- Required-No (0) to establish the connection even if its associated private-route profile is not found.
- Required-Yes (1) to drop the connection if its associated private-route profile is not found.

Dependencies: The Ascend-Private-Route-Required value overrides the setting of Private-Route-Profile-Required in the Answer-Defaults profile.

See Also: Ascend-Private-Route-Table-ID

Ascend-Cache-Refresh (56)

Description: Specifies whether successive references to a cached private-route profile reset its cache timer.

Usage: Specify one of the following values:

- Refresh-No (0) to leave the timer alone when a cached private-route table is referenced (when a session becomes active that uses the table).
- Refresh-Yes (1) to reset the cache timer each time a cached private-route table is referenced.

See Also: Ascend-Cache-Time

Ascend-Cache-Time (57)

Description: Indicates the time (in minutes) during which a private-route profile remains cached.

Usage: Specify an integer. If you do not specify the Ascend-Cache-Time attribute in a private-route profile, the profile is cached for the amount of time specified by the Default-Prt-Cache-Time parameter in the IP-Global profile.

See Also: Ascend-Cache-Refresh

Command reference entry for private-route cache management

The PrtCache command displays statistics about each cached RADIUS private-route profile, and enables you to flush profiles from the cache.

PrtCache

Description: Displays statistics about cached RADIUS private-route profiles, and enables you to flush the cache.

Permission level: Diagnostic or Update

Usage: `prtcache -s [profilename] | -f [-f] | -t`

Option	Description
<code>-s [profilename]</code>	If <i>profilename</i> is not specified, the command display statistics for all cached private-route profiles. If it is specified, the command displays statistics only for the specified private-route profile
<code>-f [-f]</code>	Flush all cached entries. The second <code>-f</code> flag specifies that all cached routes are flushed without waiting for confirmation.

Option	Description
-t	Toggle debug output.

Example: prtcache -s

Profile Name	Created	Exp After(min)	Use Count	Refresh Cache
check	12:32:53	1	0	Yes
my-route	10:32:53	23	8	No

admin> **prtcache -s check**

Profile Name	Created	Exp After(min)	Use Count	Refresh Cache
check	12:32:53	1	0	Yes

Following is a description of each field in the output:

Output field	Description
Profile Name	Name of the cached profile.
Created	Time at which the profile was created.
Exp After	Number of minutes after which the profile is removed from the cache.
Use Count	Number of times the cached profile was referred to in the past.
Refresh Cache	Specifies whether the profile's cache time is refreshed if the profile is used.

admin> **prtcache -f**

```
Flush all cached Private Route Table Profiles ? [y/n] y
All cached Private Route Table Profiles flushed.
```

If no profiles have been cached, using the -f option displays the following output:

```
admin> prtcache -f
Flush all cached Private Route Table Profiles ? [y/n] y
No cached Profiles to flush.
```

If the user does not have the required permission:

```
admin> prtcache -f
error: Command requires 'diagnose' or 'update' privileges
```

Note: All cached RADIUS private-route profiles are read only. You can delete a single cached profile by using the Delete command. To delete all cached profiles, use the PrtCache command.

Port redirection

Port redirection enables you to configure a Connection or RADIUS profile to redirect certain packet types to a specified server. An example use of this feature is to redirect Hypertext Transfer Protocol (HTTP) traffic to a Web cache server on a local network. However, port

redirection is not limited to HTTP traffic. You can use the feature to redirect any TCP or UDP packet on the basis of its protocol and port information.

Overview of Connection profile settings

Following are the relevant parameters, shown with default settings:

```
[in CONNECTION/" ":port-redirect-options]
protocol = none
port-number = 0
redirect-address = 0.0.0.0
```

Parameter	Specifies
Protocol	Protocol type. Valid settings are <code>none</code> (the default, which disables port redirection), <code>udp</code> , and <code>tcp</code> . This setting together with the Port-Number setting (next) defines a type of packet. For example, TCP with a Port-Number of 21 represents FTP traffic, and TCP with a Port-Number of 23 represents Telnet traffic. For HTTP traffic, set the parameter to <code>tcp</code> .
Port-Number	Port number to be compared with the destination port of a packet. TCP and UDP port numbers are typically assigned to services. For example, HTTP traffic uses TCP port 80. For a list of assigned port numbers, see RFC 1700, <i>Assigned Numbers</i> .
Redirect-Address	IP address to which matching packets are redirected.

Overview of RADIUS settings

RADIUS uses the following attribute-value pairs for port redirection:

RADIUS attribute	Value
Ascend-Port-Redir-Protocol (82)	Protocol type. Valid values are <code>Ascend-Proto-TCP</code> (6) and <code>Ascend-Proto-UDP</code> (17). This setting together with the <code>Ascend-Port-Redir-Portnum</code> setting (next) defines a type of packet. For example, <code>Ascend-Proto-TCP</code> with an <code>Ascend-Port-Redir-Portnum</code> setting of 21 represents FTP traffic, and <code>Ascend-Proto-TCP</code> with an <code>Ascend-Port-Redir-Portnum</code> of 23 represents Telnet traffic. For HTTP traffic, specify <code>Ascend-Proto-TCP</code> (6).
Ascend-Port-Redir-Portnum (83)	Port number to be compared with the destination port of a packet. TCP and UDP port numbers are typically assigned to services. For example, HTTP traffic uses TCP port 80. For a list of assigned port numbers, see RFC 1700, <i>Assigned Numbers</i> .
Ascend-Port-Redir-Server (84)	IP address to which matching packets are redirected.

To use these attributes, the RADIUS server must support vendor-specific attributes (VSAs) and the MAX TNT must be configured in VSA compatibility mode. Following are the relevant settings:

```
[in EXTERNAL-AUTH]
auth-type = radius
```


```
[in EXTERNAL-AUTH:rad-auth-client]
auth-radius-compat = vendor-specific
```

For details about these settings, see the *MAX TNT Reference Guide*.

Example of configuring port redirection

In this example, the MAX TNT redirects a PPP client's browser requests to a Web cache server at 1.1.1.1. The Web cache server can respond directly if a cached entry is found, or forward the browser request to its original destination if no cache entry is found. The example setup is shown in Figure 13.

Figure 13. Port redirection to an HTTP server



The following commands configure a local profile for the PPP client, redirecting its HTTP traffic to the server at 1.1.1.1:

```
admin> new connection atcp50
CONNECTION/atcp50 read
admin> set active = yes
admin> set ip-options remote-address = 2.2.2.2/32
admin> set ppp-options recv-password = test
admin> set port-redirect-options protocol = tcp
admin> set port-redirect-options port-number = 80
admin> set port-redirect-options redirect-address = 1.1.1.1
admin> write
CONNECTION/atcp50 written
```

Following is a comparable RADIUS profile:

```
atcp50 Password = "test"
  Service-Type = Framed,
  Framed-Protocol = MPP,
  Framed-IP-Address = 2.2.2.2,
  Framed-IP-Netmask = 255.255.255.255,
  Ascend-Port-Redir-Protocol = Ascend-Proto-TCP,
  Ascend-Port-Redir-Portnum = 80,
  Ascend-Port-Redir-Server = 1.1.1.1
```

IP pool chaining

An IP address pool is a range of contiguous addresses on a local IP network or subnet. The MAX TNT assigns a pool address to a caller that requests an address, and when the call terminates, the system returns the address to the pool.

```
pools-JFAN-TNT Password = "ascend"
  Service-Type = Outbound,
  Ascend-IP-Pool-Definition = "1 11.168.6.10 10",
  Ascend-IP-Pool-Definition = "2 12.168.6.10 10",
  Ascend-IP-Pool-Definition = "3 13.168.6.10 10",
  Ascend-IP-Pool-Definition = "7 17.168.6.10 10",
  Ascend-IP-Pool-Definition = "8 18.168.6.10 10",
  Ascend-IP-Pool-Definition = "9 19.168.6.10 10"
```

When IP pool chaining is enabled, contiguous pools are treated as one *pool space* with shared addresses. When the system assigns an address to an end user, it begins searching for an available address in the first pool of the chain and stops when it either finds an available address or encounters a null pool definition. So, the pools within a chain must be defined in a contiguous sequence. For example, the following profile contains two IP pool chains (pools 1, 2, 3, and pools 7, 8, 9), with each pool chain containing 30 addresses:

```
pools-JFAN-TNT Password = "ascend", Service-Type = Outbound
Ascend-IP-Pool-Chaining = IP-Pool-Chaining-Yes,
Ascend-IP-Pool-Definition = "1 11.168.6.10 10",
Ascend-IP-Pool-Definition = "2 12.168.6.10 10",
Ascend-IP-Pool-Definition = "3 13.168.6.10 10",
Ascend-IP-Pool-Definition = "7 17.168.6.10 10",
Ascend-IP-Pool-Definition = "8 18.168.6.10 10",
Ascend-IP-Pool-Definition = "9 19.168.6.10 10"
```

IP pool chaining is supported both for RADIUS-defined address pools and pools defined locally in the IP-Global profile. For example, the following settings in the IP-Global profile enable pool chaining and define a pool chain (pools 1 and 2) that contains 252 addresses:

```
[in IP-GLOBAL]
pool-chaining = yes
pool-base-address = [ 172.20.31.1 172.20.33.1 0.0.0.0 153.37.21.1 0.0+
assign-count = [ 126 126 0 30 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 +
```

Pool chaining in local profiles

Overview of local profile settings

```
[in IP-GLOBAL]
pool-chaining = no
```

```
pool-base-address = [ 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0+
assign-count = [ 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0 0+
[in CONNECTION/"":ip-options]
address-pool = 0
```

Parameter	Specifies
Pool-Chaining	Enable/disable IP pool chaining. If set to yes, the system treats contiguous IP address pools as a single extended pool space when searching for an available address to assign to a caller.
Pool-Base-Address	An array of up to 128 IP addresses to be used as the first address in a pool. These values are used with the Assign-Count values to define address pools locally. A pool chain contains all of the pools defined in sequence within the array, such as 1, 2, 3. To end a pool chain, leave a null value in the array.
Assign-Count	An array of up to 128 numbers that specify the number of addresses in a pool that starts with the corresponding Pool-Base-Address.
Address-Pool	Number of an address pool from which to acquire an address. When pool chaining is enabled, a pool number within a chain includes addresses defined in all other pools within the chain. For example, if pools 1, 2, and 3 are in a pool chain, setting this parameter to 1 has the same effect as setting it to 2 or 3.

Example of local pool chain definition

The following commands define five address pools, which form two pool chains. Notice that the pool numbers (their indexes in the Pool-Base-Address and Assign-Count arrays) are contiguous within a chain.

```
admin> read ip-global
IP-GLOBAL read

admin> set pool-chaining = yes

admin> set pool-base-address 1 = 10.1.1.1
admin> set pool-base-address 2 = 11.1.1.1
admin> set pool-base-address 3 = 12.1.1.1
admin> set assign-count 1 = 50
admin> set assign-count 2 = 50
admin> set assign-count 3 = 50
admin> set pool-base-address 7 = 13.1.1.1
admin> set pool-base-address 8 = 14.1.1.1
admin> set assign-count 7 = 50
admin> set assign-count 8 = 50

admin> write
IP-GLOBAL written
```

The following commands enable dynamic address assignment system-wide:

```
admin> read answer
ANSWER-DEFAULTS read
```

```
admin> set ip-answer assign = yes
```

```
admin> write
```

```
ANSWER-DEFAULTS written
```

The following commands configure profiles to acquire an address from the first pool chain. When the end users dial in, they can acquire an address from 10.1.1.1 to 10.1.1.51, from 11.1.1.1 to 11.1.1.51, or from 12.1.1.1 to 12.1.1.51. If no addresses are available within those ranges, the connection is refused.

```
admin> new conn jfan
```

```
CONNECTION/jfan read
```

```
admin> set active = yes
```

```
admin> set encapsulation-protocol = ppp
```

```
admin> set ppp-options recv-password = localpw
```

```
admin> set ip-options address-pool = 2
```

```
admin> write
```

```
CONNECTION/jfan written
```

```
admin> new conn ravi
```

```
CONNECTION/ravi read
```

```
admin> set active = yes
```

```
admin> set encapsulation-protocol = ppp
```

```
admin> set ppp-options recv-password = localpw
```

```
admin> set ip-options address-pool = 1
```

```
admin> write
```

```
CONNECTION/ravi written
```

Following are comparable RADIUS profiles:

```
jfan Password = "localpw"
  Service-Type = Framed-User,
  Framed-Protocol = PPP,
  Ascend-Assign-IP-Pool = 2
```

```
ravi Password = "localpw"
  Service-Type = Framed-User,
  Framed-Protocol = PPP,
  Ascend-Assign-IP-Pool = 1
```

Pool chaining in RADIUS

Whether pool chains are defined locally or in a RADIUS pool's pseudo-user profile, the pool addresses are available for dynamic assignment regardless of where the caller's profile is authenticated.

Overview of RADIUS profile settings

RADIUS servers use the following attribute-value pairs to define and apply pool chains:

RADIUS attribute	Value
Ascend-IP-Pool-Chaining (85)	<p>Enable/disable IP pool chaining in a pseudo-user profile that defines address pools. If set to IP-Pool-Chaining-Yes (1), the system treats contiguous IP address pools as a single extended pool space when searching for an available address to assign to a caller. If set to IP-Pool-Chaining-No (0), the system treats each address pool as a separate space.</p> <p>Note: When this attribute is specified in a RADIUS profile, its value overrides the Pool-Chaining setting in IP-Global profile.</p>
Ascend-IP-Pool-Definition (217)	<p>Address pool definition in a pseudo-user profile. The value has the following syntax:</p> <pre>pool-number base-addr assign-count</pre> <p>The <i>pool-number</i> value is an integer that identifies the pool. A pool chain contains all of the pools defined in sequence, such as 1, 2, 3. To end a pool chain, leave a gap in the sequence of <i>pool-number</i> values. The <i>base-addr</i> value is an IP address to be used as the first address in a pool, and the <i>assign-count</i> value specifies the number of addresses in a pool.</p>
Ascend-Assign-IP-Pool (218)	<p>Number of the address pool from which the RADIUS user profile should acquire an address. When pool chaining is enabled, a pool number within a chain includes addresses defined in all other pools within the chain. For example, if pools 1, 2, and 3 are in a pool chain, setting this value to 1 has the same effect as setting it to 2 or 3.</p>

To use these attributes, the RADIUS server must support vendor-specific attributes (VSAs) and the MAX TNT must be configured in VSA compatibility mode. Following are the relevant settings:

```
[in EXTERNAL-AUTH]
auth-type = radius

[in EXTERNAL-AUTH:rad-auth-client]
auth-radius-compatible = vendor-specific
```

For details about these settings, see the *MAX TNT Reference Guide*.

Example of pool chaining in RADIUS

The following pseudo-user profile defines five address pools, which form two pool chains. Notice that the pool numbers are contiguous within a chain.

```
pools-JFAN-TNT Password = "ascend"
  Service-Type = Outbound,
  Ascend-IP-Pool-Chaining = IP-Pool-Chaining-Yes,
  Ascend-IP-Pool-Definition = "1 10.1.1.1 50",
  Ascend-IP-Pool-Definition = "2 11.1.1.1 50",
  Ascend-IP-Pool-Definition = "3 12.1.1.1 50"
  Ascend-IP-Pool-Definition = "7 13.1.1.1 50",
  Ascend-IP-Pool-Definition = "8 14.1.1.1 50"
```

The following commands configure local Connection profiles to acquire an address from the first pool chain. When the end users dial in, they can acquire an address from 13.1.1.1 to 13.1.1.51, or from 14.1.1.1 to 14.1.1.51. If no addresses are available within those ranges, the connection is refused.

```
admin> new conn hanif
CONNECTION/hanif read

admin> set active = yes

admin> set encapsulation-protocol = ppp

admin> set ppp-options recv-password = localpw

admin> set ip-options address-pool = 7

admin> write
CONNECTION/hanif written

admin> new conn hasnain
CONNECTION/hasnain read

admin> set active = yes

admin> set encapsulation-protocol = ppp

admin> set ppp-options recv-password = localpw

admin> set ip-options address-pool = 8

admin> write
CONNECTION/hasnain written
```

Following are comparable RADIUS user profiles:

```
hanif Password = "localpw"
    Service-Type = Framed-User,
    Framed-Protocol = PPP,
    Ascend-Assign-IP-Pool = 7

hasnain Password = "localpw"
    Service-Type = Framed-User,
    Framed-Protocol = PPP,
    Ascend-Assign-IP-Pool = 8
```

Parameter reference entry

Pool-Chaining

Description: Enable/disable IP pool chaining. If set to *yes*, the system treats contiguous IP address pools as a single extended pool space when searching for an available address to assign to a caller.

Usage: Specify *yes* or *no*. The default is *no*.

Example: `set pool-chaining = yes`

Dependencies: Address pools must be defined, either locally or in RADIUS pseudo-user profiles. Address assignment must be enabled in the Answer-Defaults profile.

Location: IP-Global

See Also: Pool-Base-Address, Assign-Count, Address-Pool, Assign-Address

RADIUS attribute reference entry

Ascend-IP-Pool-Chaining (85)

Description: Enable/disable IP pool chaining in a pseudo-user profile that defines address pools. If set to IP-Pool-Chaining-Yes (1), the system treats contiguous IP address pools as a single extended pool space when searching for an available address to assign to a caller. If set to IP-Pool-Chaining-No (0), the system treats each address pool as a separate space.

Usage: Specify IP-Pool-Chaining-Yes (1) to enable IP pool chaining or IP-Pool-Chaining-No (0) to disable it.

Example: `Ascend-IP-Pool-Chaining = IP-Pool-Chaining-Yes`

Dependencies: Address pools must be defined, either locally or in RADIUS pseudo-user profiles. Address assignment must be enabled in the Answer-Defaults profile. VSA-compatibility mode must be enabled in the External-Auth profile. When this attribute is specified in a RADIUS profile, its value overrides the setting in IP-Global profile.

Location: RADIUS pools pseudo-user profiles

See Also: Ascend-IP-Pool-Definition (217), Ascend-Assign-IP-Pool (218)

Don't Fragment option added to Ping command

The Ping command now supports the `-f` option to set the Don't Fragment (DF) bit in the IP header of Ping packets. The usage is as follows:

```
ping -f host
```

For example:

```
admin> ping -f 10.1.1.1
```

Commands for monitoring multipath route caching

Multipath route caching enables the IP cache entries on slot card interfaces to distribute the load of packets meant for a particular destination among multiple gateways that have been configured to that destination.

New MPRT debug-level command

A new debug-level MPRT command is now available both on the shelf controller when debug permissions are enabled, and on individual slot cards. On the shelf controller, the command has the following usage:

```
admin> mprrt ?  
usage: mprrt [-l|-t]
```

The `-l` option produces long-format output, which includes gateway information for each next hop of multipath routes. The `-t` option toggles display of MPRT debug messages.

On a slot card, the command also supports the `-e` option for enabling (the default) and the `-d` option for disabling multipath route caching:

Built-in features in MAX TNT TAOS 8.0.0

IP routing

```
ether2-1/5> mprt ?
usage: mprt [-l|-e|-d|-t]
```

By default, the command displays destinations for which multipath routes have been configured, and the number of next hops for each one. For example:

```
admin> mprt
MP Route
0.0.0.0/0
                                Num Paths = 3
20.0.0.0/8
                                Num Paths = 3
194.194.194.0/24
                                Num Paths = 3
```

With the `-l` option, the MPRT command displays shelf and IP address information for the next hops. For example:

```
ether2-1/5> mprt -l
MP Route      Gateway      Shelf/Slot IF Addr      Mtu Switched
0.0.0.0/0
              10.1.2.13   ( 1/14)    10.1.2.38    1524  0
              10.1.2.9    ( 1/16)    10.1.2.34    1524  0
              10.1.2.5    ( 1/15)    10.1.2.42    1524  0
              Num Paths = 3
20.0.0.0/8
              10.1.2.13   ( 1/14)    10.1.2.38    1524  0
              10.1.2.9    ( 1/16)    10.1.2.34    1524  0
              10.1.2.5    ( 1/15)    10.1.2.42    1524  0
              Num Paths = 3
194.194.194.0/24
              10.1.2.13   ( 1/14)    10.1.2.38    1524  0
              10.1.2.9    ( 1/16)    10.1.2.34    1524  0
              10.1.2.5    ( 1/15)    10.1.2.42    1524  0
              Num Paths = 3
```

Changes to the IPCache Cache command

IPCache Cache command output now shows route type and multipath information. Like the MPRT command, this command is supported both on the shelf controller and on the slot cards. The following example shows command output on the shelf controller:

```
admin> ipcache cache
Hsh    Address      Gateway      Ifname    Sh/Sl/T    MTU
20     50.0.0.20        10.168.26.74 wan392    1/14/D     1524
40     20.0.0.40        20.0.0.40    ie1-3-1   1/3 /S     1500
```

```
Cache Limit 0 Cache Count 2 Cache over limit 0 No.packets 9
```

```
Mem Usage: Allocated 1k bytes
Free block count 22
```

The following example shows command output on a slot card:

```
admin> open 1 3

ether2-1/3> ipcache cache
Hsh Address      Gateway      Sh/Sl/T Switched  MTU    MPath
0   99.1.1.1      10.168.21.30 1/14/D  0        1524   Y/0.0.0.0/0
```

20	50.0.0.20	10.168.28.170	1/15/D	85068	1524	Y/0.0.0.0/0
40	20.0.0.40	20.0.0.40	1/3 /S	0	1500	N

The new T (Type) column following the shelf and slot numbers can specify D for dynamic cache entries or S for static cache entries. The MPath column indicates whether the cache entry is derived from multipath routes. If it represents a multipath route, the column indicates Y and the destination address. If it is not a multipath route, the column indicates N.

Store-and-forward IP fax

The store-and-forward IP fax feature enables the MAX TNT to interact with a third-party fax server, such as the servers provided by Open Port Technology, Inc. Fax-over-IP technology enables ISPs and corporate hubs to use the Internet to deliver faxes.

When the IP fax feature is enabled in the MAX TNT, the system acts as a remote access server (RAS), accepting fax calls using the same ports and telephone lines used for dial-in modem connections. The unit also performs modem dial-out functions to deliver faxes from the Internet to fax machines on the Public Switched Telephone Network (PSTN).

Incoming and outgoing IP faxes

Figure 14 shows the basic structure of an incoming IP fax operation. The MAX TNT receives an *incoming fax* from the PSTN and interacts with the fax server to transfer it to the Internet. The transfer to the Internet is transparent to the person sending a fax, because a hardware device called a *redialer* is connected to the fax machine. The redialer intercepts the number dialed on the fax machine and initiates a call to the MAX TNT instead. When the fax server begins transferring the fax to the Internet, the redialer and the MAX TNT become transparent pipes for the fax data.

Figure 14. Incoming IP fax from fax machine to Internet





Figure 15 shows the basic structure of an outgoing IP fax operation. The fax server receives an *outgoing fax* from the Internet and interacts with the MAX TNT to transfer it to the PSTN. The fax server logs into the MAX TNT and is authenticated before seizing one of the unit's modems for dial-out to the destination fax machine.

Figure 15. Outgoing IP fax from Internet to fax machine



System parameters for IP fax modem usage

To send faxes, the fax server logs into the MAX TNT, gains control of one of its modems, and dials out. The fax server configuration specifies the IP address of the MAX TNT and (optionally) one or more trunk groups for IP fax use. In addition to the IP fax login and port parameters that enable the fax server to log in, which are described in the next section, the following parameters in the System profile affect the resources available for outgoing fax calls. The parameters are shown here with their default settings:

```

[ in SYSTEM ]
use-trunk-groups = no
num-digits-trunk-groups = 1
parallel-dialing = 2

[ in T1/{ any-shelf any-slot 0 }:line-interface ]
default-call-type = digital

[ in T1/{ any-shelf any-slot 0 }:line-interface:channel-config[1] ]
trunk-group = 9

[ in E1/{ any-shelf any-slot 0 }:line-interface ]
default-call-type = digital

[ in E1/{ any-shelf any-slot 0 }:line-interface:channel-config[1] ]
trunk-group = 9

[ in SWAN { any-shelf any-slot 0 }:line-config ]
trunk-group = 9

[ in CALL-ROUTE { { { any-shelf any-slot 0 } 0 } 0 } ]
trunk-group = 0

```

Parameter	Specifies
Use-Trunk-Groups	Enable/disable the use of trunk groups in the MAX TNT. If set to no (the default), the Num-Digits-Trunk-Groups and Trunk-Group settings do not apply. If set to yes, all channels must be assigned trunk group numbers.
Num-Digits-Trunk-Groups	Number of digits to allow for trunk groups. Currently, the IP fax server supports 2-digit trunk groups, but the trunk group number specification must be within the range of 2 to 9. The MAX TNT must agree with the fax server about the number of digits in a trunk group number, or telephone numbers are not parsed correctly and calls fail. For details, see “Support for multidigit trunk groups” on page 37.
Parallel-Dialing	Total number of dial-out calls that the MAX TNT can place at the same time.

Parameter	Specifies
Default-Call-Type	Default call type for calls on non-ISDN T1 or E1 lines. This parameter must be set to <code>voice</code> for IP fax over inband signaling.
Trunk-Group	Trunk group number (from 2 to 9). For a network line, this parameter assigns channels to a trunk group. In a Call-Route profile, it specifies that calls received on that trunk group will be routed to the shelf, slot, and port specified in the profile's index.

Assigning bandwidth for typical IP fax usage

After the fax server has control of a digital modem, it dials the call on any available channel unless the fax server configuration specifies a trunk group number. In this case, the fax server uses an available channel within that trunk group. If no channels in that trunk group are available, the MAX TNT returns a Trunk Group Not Available code to the fax server, which tries the call again later.

For example, the following commands configure the system to use 2-digit trunk groups, and configure a T1 line in trunk group 5. (Fewer than 24 channels can be assigned to a trunk group if appropriate.) If the fax server configuration also specifies 2-digit trunk groups and trunk group 5, these channels are available for IP fax usage.

```
admin> read system
SYSTEM read

admin> set use-trunk-groups = yes

admin> set num-digits-trunk-groups = 2

admin> write
SYSTEM read

admin> read t1 { 1 5 7 }
T1/{ shelf-1 slot-5 7 } read

admin> set line default-call-type = voice

admin> set line channel 1 trunk = 5
admin> set line channel 2 trunk = 5
admin> set line channel 3 trunk = 5
admin> set line channel 4 trunk = 5
admin> set line channel 5 trunk = 5
admin> set line channel 6 trunk = 5
admin> set line channel 7 trunk = 5
admin> set line channel 8 trunk = 5
admin> set line channel 9 trunk = 5
admin> set line channel 10 trunk = 5
admin> set line channel 11 trunk = 5
admin> set line channel 12 trunk = 5
admin> set line channel 13 trunk = 5
admin> set line channel 14 trunk = 5
```

Built-in features in MAX TNT TAOS 8.0.0

Store-and-forward IP fax

```
admin> set line channel 15 trunk = 5
admin> set line channel 16 trunk = 5
admin> set line channel 17 trunk = 5
admin> set line channel 18 trunk = 5
admin> set line channel 19 trunk = 5
admin> set line channel 20 trunk = 5
admin> set line channel 21 trunk = 5
admin> set line channel 22 trunk = 5
admin> set line channel 23 trunk = 5
admin> set line channel 24 trunk = 5
admin> write
T1/{ shelf-1 slot-5 7 } written
```

Configuring a typical Call-Route profile

After assigning the trunk group, you must create a Call-Route profile to direct outbound calls that specify trunk group 5 to the line. For example:

```
admin> new call-route { { { shelf-1 slot-5 7 } 0 } 0 }
CALL-ROUTE/{ { { shelf-1 slot-5 7 } 0 } 0 } read
admin> set trunk-group = 5
admin> set call-route-type = trunk-call
admin> write
CALL-ROUTE/{ { { shelf-1 slot-5 7 } 0 } 0 } written
```

Specifying the maximum number of parallel dial-outs

The Parallel-Dialing parameter limits the number of dial-out calls that the system can place at one time. If the maximum number of dial-out calls is being processed and a dial-out request is made, the system queues the request and processes it at the earliest possible opportunity.

This operation is transparent to the fax server, except that the modems can time out if a dial-out request is delayed more than 30 to 40 seconds. Following is an example that sets Parallel-Dialing to the maximum value for T1:

```
admin> read system
SYSTEM read
admin> set parallel-dialing = 64
admin> write
SYSTEM read
```

Configuring the MAX TNT for IP fax

Following are the IP fax parameters that enable the MAX TNT to interact with a third-party fax server. The parameters are shown with their default settings:

```
[in IP-FAX]
ip-fax-enabled = no
outgoing-fax-port = 10001
server-login = ""
```

```
server-password = ""
incoming-fax-port = 0
all-calls-are-fax = no

[in IP-FAX:fax-dnis]
fax-dnis[1] = ""
fax-dnis[2] = ""
fax-dnis[3] = ""
fax-dnis[4] = ""
fax-dnis[5] = ""
fax-dnis[6] = ""
fax-dnis[7] = ""
fax-dnis[8] = ""

[in IP-FAX:fax-servers]
fax-servers[1] = 0.0.0.0
fax-servers[2] = 0.0.0.0
fax-servers[3] = 0.0.0.0
fax-servers[4] = 0.0.0.0
fax-servers[5] = 0.0.0.0
```


Parameter	Specifies
IP-Fax-Enabled	Enable/disable IP fax support in the MAX TNT. It is disabled by default.
Outgoing-Fax-Port	TCP port on which to accept outgoing fax data from a fax server. (Outgoing fax data is received from the Internet and requires a dialout to a destination fax machine.) The default is 10001.
Server-Login Server-Password	Name and password used to authenticate the fax server as part of an outgoing fax session. When the fax server receives a fax from the Internet, it connects to the MAX TNT and sends a name and password. The MAX TNT compares the values sent to the Server-Login and Server-Password settings.
Incoming-Fax-Port	TCP port on which the fax server listens for incoming fax data. (Incoming fax data is received from a fax machine redialer.) The default is zero.
All-Calls-Are-Fax	Enable/disable the handling of all incoming calls as IP fax calls. When this parameter is set to <code>no</code> (the default), the MAX TNT recognizes incoming fax calls by matching the caller's DNIS number to one of the configured Fax-DNIS numbers (next). If set to <code>yes</code> , IP fax service can be supported where DNIS is not available.
Fax-DNIS [1–8]	Up to 8 DNIS numbers. The MAX TNT compares the DNIS number supplied in the PRI setup message of an incoming call to the configured numbers. If the match is not exact, the unit does not start the IP fax functionality.

Parameter	Specifies
Fax-Servers [1–5]	<p>IP address of up to 5 fax servers. The fax server systems are typically on the local IP network, but local connectivity is not a requirement.</p> <p>The MAX TNT first tries to connect to the fax server at the first specified address. If the unit receives no response, it tries to connect to the second address. If the unit still receives no response, it tries the third, and so forth. Once the MAX TNT connects to a fax server successfully, it continues to use that address for subsequent connections until a connection attempt fails, at which point it tries the next configured address.</p>

Example of an IP fax configuration for incoming faxes

Figure 16 shows a MAX TNT receiving an incoming fax across the PSTN. The unit then initiates a TCP session with a fax server, which authenticates the incoming call. (The fax server might use RADIUS, as shown in Figure 16, or a method proprietary to that server.) If the fax server authenticates the call successfully, it dials out to the remote fax server on one of the MAX TNT modems. When the fax transmission has completed, the fax server terminates the TCP session and the MAX TNT regains control of its modem.

Figure 16. Receiving and forwarding incoming IP faxes



Following is an example of an IP fax setup that enables the MAX TNT to handle incoming fax calls as shown in Figure 16:

```
admin> new ip-fax
IP-FAX read

admin> set ip-fax-enabled = yes
admin> set incoming-fax-port = 1234
admin> set fax-dnis 1 = 2222
admin> set fax-servers 1 = 10.1.2.34
admin> set fax-servers 2 = 10.1.2.56
```



```
admin> list
ip-fax-enabled = yes
outgoing-fax-port = 10001
server-login = ""
server-password = ""
incoming-fax-port = 1234
all-calls-are-fax = no
fax-dnis = [ 2222 "" "" "" "" "" "" "" ]
fax-servers = [ 10.1.2.34 10.1.2.56 0.0.0.0 0.0.0.0 0.0.0.0 ]

admin> write
IP-FAX written
```

With this configuration, the IP fax is processed as follows:


- 1 An end user sends a fax to 123-555-1111.
- 2 The sending fax machine receives a dial tone from the redialer (which is directly connected to the fax machine) and dials 123-555-1111.
- 3 The redialer intercepts the call, stores the destination telephone number, and dials its configured number for the MAX TNT (456-555-2222).
- 4 The MAX TNT receives the call and identifies it as a fax call by comparing the call's DNIS number to the Fax-DNIS values in the IP-Fax profile.
- 5 If the DNIS numbers match (or if the unit is configured to treat all incoming calls as IP fax calls), the MAX TNT generates an answer tone at 400 Hz to initiate dual-tone multifrequency (DTMF) communication with the redialer. Then the unit decodes the incoming DTMF sequence from the redialer, which contains the account number of the redialer and the destination telephone number 123-555-1111.
- 6 The MAX TNT initiates a connection to the fax server, sending the caller's account number and destination telephone number in the first TCP packet.
- 7 If the fax server authenticates the call successfully using this information, the MAX TNT answers the incoming fax call. If authentication fails, the connection is cleared.
- 8 Following successful authentication, the MAX TNT and fax server establish a TCP session, and the MAX TNT transfers control of an available modem to the fax server for the incoming call. If no send or receive activity occurs for more than 2 minutes, the session is terminated and resources are freed.

Note: For fax accounting, a fax session starts when a modem resource is allocated and stops when a session is terminated.

Example of an IP fax configuration for outgoing faxes

Figure 17 shows a MAX TNT forwarding a fax received by the fax server from the Internet. The fax server logs into the MAX TNT using the configured Server-Login and Server-Password, and initiates a modem dial-out session to forward the fax on the PSTN. When the fax transmission has completed, the fax server terminates the TCP session and the MAX TNT regains control of its modem.

Figure 17. Sending an outgoing IP fax to a fax machine



Following is an example of an IP fax setup that enables the MAX TNT to handle outgoing fax calls as shown in Figure 17:

```
admin> new ip-fax
IP-FAX read

admin> set ip-fax-enabled = yes

admin> set server-login = ipfax

admin> set server-password = works

admin> list
ip-fax-enabled = yes
outgoing-fax-port = 10001
server-login = ipfax
server-password = works
incoming-fax-port = 0
All calls are Fax = no
fax-dnis = [ "" "" "" "" "" "" "" "" ]
fax-servers = [ 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 0.0.0.0 ]

admin> write
IP-FAX written
```

With this configuration, the MAX TNT processes an IP fax as follows:

- 1 The fax server on the local network receives fax data across the Internet from a remote fax server.
- 2 The fax server initiates a connection to the MAX TNT, sending its login name and password in the first TCP packet.
- 3 If the login name and password match the Server-Login and Server-Password in the IP-Fax profile, the MAX TNT establishes a TCP session with the fax server. If authentication fails, the connection is cleared.
- 4 After authentication, the MAX TNT transfers control of an available modem to the fax server.
- 5 The fax server sends modem commands encapsulated in TCP packets, initiates a connection to the destination fax machine, and sends the spooled data. If no send or receive activity occurs for more than 2 minutes, the session is terminated and resources are freed.

Note: For fax accounting, a fax session starts when a modem resource is allocated and stops when a session is terminated.

Fax hangup codes and disconnect cause codes

Conexant supplies two fax hangup codes:

- +FHNG 1: when fax tones are recognized but the handshake fails
- +FHNG 11: when no fax tones are recognized at the far end

ISDN disconnect cause codes are returned when fax calls fail, if they are available as part of the fax hangup codes. To avoid conflict with codes returned by modems and with codes returned by other units, the fax cause codes add 1000 to the standard codes so that they are in the range of 1000 - 1255. For example, Far End Busy (ISDN Code 17) is returned as +FHNG 1017, and Far End Did Not Answer (go off-hook) is returned as +FHNG 1018.

IP fax call accounting

In previous releases, more accounting information was available for an incoming call than for an outgoing call. Because the IP fax feature creates a large volume of outgoing calls, SNMP, RADIUS, and Syslog call-accounting information has been expanded to include the following additional accounting information for outgoing IP fax calls:

- A call-connected timestamp, showing the length of the call
- The trunk group used for particular channels on an outgoing call
- The destination telephone number dialed from the MAX TNT
- The shelf, slot, line and channel number at which the call originates
- The total bytes sent and received (in SNMP and RADIUS only)
- The transmit and receive baud rate (in SNMP and RADIUS only)

Note: For accounting purposes, a fax session starts when a modem resource is allocated and stops when the session is terminated.

SNMP changes for IP fax operation

SNMP provides additional call information in the following fields:

MIB field name	Reports
eventCurrentService: ipFax (19)	Service ipFax is available for an IP fax call when the event type is callOriginated(1).
eventTrunkGroup (24)	Trunk group used for outgoing calls only. This information is available when the event type is callCleared (9).
eventCalledPartyID	Telephone number dialed for an outgoing call. Currently, the eventCalledPartyID is equivalent to the DNIS Dialed Number ID for an incoming call. On the outgoing call, this field represents the telephone number dialed. This information is available when the event type is callCleared (9).
eventSlotNumber	Slot number at which the call originated. This information is available when the event type is callCleared(3).

MIB field name	Reports
eventSlotLineNumber	Line at which the call originated. This information is available when the event type is callCleared(3).
eventSlotChannelNumber	Channel at which the call originated. This information is available when the event type is callCleared(3).
eventTimeStamp	For an IP fax call, the time that the modem is reserved for an outgoing call request; for any other type of call, this field reports the actual connected time. This information is available when the event type is callCleared(3).
eventInOctets	Total received bytes for the call. This information is available when the event type is callCleared(3).
eventOutOctets	Total transmitted bytes for the call. This information is available when the event type is callCleared(3).
eventXmitRate	Negotiated transmitted baud rate used throughout the call. This information is available when the event type is callCleared(3). For IP fax, transmitted and received baud rates are the same.
eventDataRate	Negotiated received baud rate used throughout the call. This information is available when the event type is callCleared(3). For IP fax, transmitted and received baud rates are the same.
eventUserIPAddress	User's IP address. This information is available when the event type is nameChanged(5).
eventUserName	Username. This information is available when the event type is callOriginated(1).
eventModemSlotNumber	Slot in which the modem is located. This information is available when the event type is callOriginated(1).
eventModemOnSlot	Modem in use. This information is available when the event type is callOriginated(1).
ssnActiveUserName	Active username.
ssnActiveUserIPAddress	Active user's IP address.
ssnActiveCurrentService: ipFax(19)	ipFax(19) service is in use for an outgoing IP fax call.

RADIUS changes for IP fax operation

The following RADIUS attributes, which appear in Accounting Stop packets, now reflect outgoing call values in addition to their existing incoming call values:

RADIUS attribute	Value
NAS-Port	<p>Shelf, slot, line and channel number from which the outgoing call originates. The value appears in the following binary format:</p> <p>FFSS SSSL LLLC CCCC</p> <p>FF specifies the shelf number.</p> <p>SSSS specifies the slot number.</p> <p>LLLLL specifies the line number.</p> <p>CCCC specifies the channel number.</p> <p>Each value is zero-based. For example, given the decimal number 13348, whose binary equivalent is 0011 0100 0010 0100:</p> <p>00: shelf number = 1</p> <p>1101: slot number =14</p> <p>00001: line number = 2</p> <p>0100: channel number =5</p>
Acct-Session-Time	Total connection time for a call. For an outgoing IP fax call, the time period begins when the modem is reserved, and ends when the call is terminated.
Client-Port-DNIS	Called number for an outgoing call.
Ascend-Modem-PortNo	Modem port used for the call.
Ascend-Modem-SlotNo	Number of the slot in which the modem card is physically located.
Ascend-Modem-ShelfNo	Number of the shelf on which the modem card is located.
Acct-Input-Octets	Total received bytes for the call.
Acct-Output-Octets	Total transmitted bytes for the call.
Ascend-Xmit-Rate	Negotiated transmitted baud rate for the call. For IP fax, transmitted and received baud rates are the same.
Ascend-Data-Rate	Negotiated received baud rate for the call. For IP fax, transmitted and received baud rates are the same.

In addition, the Ascend-CBCP-Trunk-Group attribute (115) has been modified to apply to outgoing IP fax calls.

Ascend-CBCP-Trunk-Group (115)

Description: Assigns the callback or outgoing IP fax call to a MAX TNT trunk group. The value in Ascend-CBCP-Trunk-Group is prepended to the number that the MAX TNT dials for callback or outgoing fax.

Usage: Specify a trunk group number from 1 to 9.

Dependencies: Ascend-CBCP-Trunk-Group applies only if one or both of the following conditions are true:

- Callback Control Protocol (CBCP) is negotiated for a connection.
- The call is an outgoing IP fax call and trunk groups are enabled in the System profile.

Syslog changes for IP fax operation

The following Syslog message reflects the time at which the modem was reserved:

```
LOG info, Shelf 1, Controller, Time: 15:36:40--  
[1/1/13/0] [MBID 13] Assigned to Port
```

The following message displays the dial-out number, trunk group, modem slot, and modem number when the call is placed:

```
LOG info, Shelf 1, Controller, Time: 15:37:07--  
[1/1/13/0] [MBID 13; ->97476799] Outgoing Call, 97476799, Trunk 8
```

When the call is connected, its shelf, slot, line, and channel are displayed in the following message:

```
LOG info, Shelf 1, Controller, Time: 15:37:13--  
[1/14/2/5] [MBID 13; ->97476799] Call Connected
```

When the call is terminated, the time, modem slot, and modem number are displayed.

```
LOG info, Shelf 1, Controller, Time: 15:38:00--  
[1/1/13/0] [MBID 13; ->97476799] Call Terminated
```

Redialer support on MultiDSP card for store-and-forward fax

When a redialer device is attached to a fax machine, it waits for a tone at 400 Hz. After receiving the tone, the redialer transmits the destination fax number to the MAX TNT as DTMF digits. With MAX TNT TAOS 8.0.0, the MultiDSP card transmits the 400-Hz tone and detects incoming DTMF digits.

Extension features in MAX TNT TAOS 8.0.0

Global Digital Access

PHS Internet Access Forum Standard (PIAFS) 2.1 on MultiDSP

Support for PIAFS version 2.1 is new in MAX TNT TAOS 8.0.0. In earlier TAOS 7.x releases, the MAX TNT supported PIAFS 1.0 (which supports a fixed data rate of 32Kbps) and PIAFS 2.0 (which supports a fixed data rate of either 32Kbps or 64Kbps for the duration of a call).

With a MultiDSP card installed, the MAX TNT supports the PIAFS protocol required for the Personal Handyphone System (PHS). PHS service is currently available only with Japan PRI signaling. With PIAFS 2.1, the MAX TNT supports a data rate that switches between 32Kbps and 64Kbps during a call, depending on what the wireless bandwidth permits. PIAFS version 2.1 has an enhanced link-level protocol that supports dynamic switching of data rates between 32 Kbps and 64 Kbps.

When the PHS-Support and PHS-2-1 licenses have been enabled, the system creates a new Call-Route profile for each installed MultiDSP card. The new Call-Route profile sets the Call-Route-Type parameter to `phs-call-type`, as shown in the following sample profile:

```
admin> get call-route { { { 1 12 0 } 0 } 2 }  
[in CALL-ROUTE/{ { { shelf-1 slot-12 } 0 2 }]  
index* = { { { shelf-1 slot-12 0 } 0 } 2 }  
trunk-group = 0  
telephone-number = ""  
preferred-source = { { any-shelf any-slot 0 } 0 }  
call-route-type = phs-call-type
```

This Call-Route-Type setting enables the system to route PHS calls to the card.

Asynchronous V.110 on MultiDSP

Following are the V.110 supported features:

- Asynchronous, answer mode (answers but does not call out), with 1 start bit, 8 data bits, and 1 stop bit.
- 48 channels per 48-port MultiDSP card, 96 channels per 96-port MultiDSP card
- Supported rates are 2400, 4800, 9600, 19200 and 38400 bits per second.

When a V.110 license has been enabled, the system creates a new Call-Route profile for each installed MultiDSP card. The new Call-Route profile sets the Call-Route-Type parameter to `v110-call-type`, as shown in the following sample profile:

```
admin> get call-route { { { 1 6 0 } 0 } 2 }  
[in CALL-ROUTE/{ { { shelf-1 slot-6 } 0 2 }]  
index* = { { { shelf-1 slot-6 0 } 0 } 2 }  
trunk-group = 0  
phone-number = ""  
preferred-source = { { any-shelf any-slot 0 } 0 }  
call-route-type = v110-call-type
```

This setting enables the system to route V.110 calls to the card.

The Callroute command lists the MultiDSP card's DSP channels as available host-side call routing entries. For example:

```
admin> callroute -ah
1:06:01/0    2 0:00:00/0    v110-call-type    0 0
1:06:03/0    2 0:00:00/0    v110-call-type    0 0
1:06:05/0    2 0:00:00/0    v110-call-type    0 0
1:06:07/0    2 0:00:00/0    v110-call-type    0 0
1:06:09/0    2 0:00:00/0    v110-call-type    0 0
1:06:11/0    2 0:00:00/0    v110-call-type    0 0
1:06:13/0    2 0:00:00/0    v110-call-type    0 0
1:06:15/0    2 0:00:00/0    v110-call-type    0 0
.
.
.
1:06:95/0    2 0:00:00/0    v110-call-type    0 0
```

R2 CLID processing for New Zealand

With MAX TNT TAOS 8.0.0, Calling-line ID (CLID) processing for R2 signaling is supported for systems in New Zealand. Following are the relevant parameter settings:

```
[in E1/{ any-shelf any-slot 0 }:line-interface]
signaling-mode = e1-new-zealand-signaling
caller-id = get-caller-id
```

To process CLID information sent by the switch, MAX TNT systems in New Zealand must use the E1-New-Zealand-Signaling setting. For example:

```
admin> read e1 {1 2 5}
E1/{ shelf-1 slot-2 5 } read
admin> set line enabled = yes
admin> set line signaling = e1-new-zealand-signaling
admin> set line switch-type = switch-cas
admin> set line group-b-signal = signal-b-6
admin> set line group-ii-signal = signal-ii-1
admin> set line caller-id = get-caller-id
admin> write
E1/{ shelf-1 slot-2 5 } written
```

For details about configuring CLID authentication in a Connection profile, see the *MAX TNT Network Configuration Guide*.

R2 CLID processing for Thailand

With MAX TNT TAOS 8.0.0, Calling-line ID (CLID) processing for R2 signaling is supported for systems in Thailand. Following are the relevant parameter settings:

```
[in E1/{ any-shelf any-slot 0 }:line-interface]
signaling-mode = e1-thailand-signaling
caller-id = get-caller-id
```

To process CLID information sent by the switch, MAX TNT systems in Thailand must use the E1-Thailand-Signaling setting. For example:

```
admin> read e1 {1 2 5}
E1/{ shelf-1 slot-2 5 } read

admin> set line enabled = yes

admin> set line signaling = e1-thailand-signaling

admin> set line switch-type = switch-cas

admin> set line group-b-signal = signal-b-6

admin> set line group-ii-signal = signal-ii-1

admin> set line caller-id = get-caller-id

admin> write
E1/{ shelf-1 slot-2 5 } written
```

For details about configuring CLID authentication in a Connection profile, see the *MAX TNT Network Configuration Guide*.

R2 CLID processing for Israel

With MAX TNT TAOS 8.0.0, Calling-line ID (CLID) processing for R2 signaling is supported for systems in Israel. Following are the relevant parameter settings:

```
[in E1/{ any-shelf any-slot 0 }:line-interface]
signaling-mode = e1-israel-signaling
caller-id = get-caller-id
```

To process CLID information sent by the switch, MAX TNT systems in Israel must use the E1-Israel-Signaling setting. For example:

```
admin> read e1 {1 2 5}
E1/{ shelf-1 slot-2 5 } read

admin> set line enabled = yes

admin> set line signaling = e1-israel-signaling

admin> set line switch-type = switch-cas

admin> set line group-b-signal = signal-b-6

admin> set line group-ii-signal = signal-ii-1

admin> set line caller-id = get-caller-id

admin> write
E1/{ shelf-1 slot-2 5 } written
```

For details about configuring CLID authentication in a Connection profile, see the *MAX TNT Network Configuration Guide*.

R2 CLID processing for Mexico

With MAX TNT TAOS 8.0.0, Calling-line ID (CLID) processing for R2 signaling is supported for systems in Mexico. Following are the relevant parameter settings:

```
[in E1/{ any-shelf any-slot 0 }:line-interface]
signaling-mode = e1-mexico-signaling
caller-id = get-caller-id
```

To process CLID information sent by the switch, MAX TNT systems in Mexico must use the E1-Mexico-Signaling setting. For example:

```
admin> read e1 {1 2 5}
E1/{ shelf-1 slot-2 5 } read

admin> set line enabled = yes

admin> set line signaling = e1-mexico-signaling

admin> set line switch-type = switch-cas

admin> set line group-b-signal = signal-b-6

admin> set line group-ii-signal = signal-ii-1

admin> set line caller-id = get-caller-id

admin> write
E1/{ shelf-1 slot-2 5 } written
```

For details about configuring CLID authentication in a Connection profile, see the *MAX TNT Network Configuration Guide*.

R2 signaling for Kuwait

R2 signaling is supported for systems in Kuwait. Following is the relevant parameter setting:

```
[in E1/{ any-shelf any-slot 0 }:line-interface]
signaling-mode = e1-kuwait-signaling
```

To use R2 signaling, MAX TNT systems in Kuwait must use the E1-Kuwait-Signaling setting. For example:

```
admin> read e1 {1 2 5}
E1/{ shelf-1 slot-2 5 } read

admin> set line enabled = yes

admin> set line signaling = e1-kuwait-signaling

admin> set line switch-type = switch-cas

admin> set line group-b-signal = signal-b-6

admin> set line group-ii-signal = signal-ii-1

admin> write
E1/{ shelf-1 slot-2 5 } written
```

Support for ISDN network-side emulation (T1 and E1)

With MAX TNT TAOS 8.0.0, you can configure PRI lines to use either network-side or user-side ISDN emulation. Previously, PRI lines on the MAX TNT supported only user-side emulation. Following is the relevant parameter, shown with its default setting:

```
[in T1/{ any-shelf any-slot0 }:line-interface]
isdn-emulation-side = te

[in E1/{ any-shelf any-slot0 }:line-interface]
isdn-emulation-side = te
```

ISDN is a nonsymmetrical protocol used by telephone carriers to provide digital services to end users. There are no ISDN links between telephone carrier Central Offices (COs). ISDN links exist only between the CO and the customer. Therefore, an ISDN link can be viewed as having two sides—the network side, or network terminating (NT) equipment, and the user side, or terminal equipment (TE). The user side can connect only to the network side, and vice

versa. Both the network side and the user side perform the same functions, but the format of the messages is different. For example, the network side must always set a bit and the user side must always clear it. These differences allow either side to determine whether the other end is the right one.

ISDN emulation enables you to build, send, receive, and process ISDN data. ISDN monitoring, on the other hand, allows you only to decode the ISDN data.

ISDN-Emulation-Side

Description: Specifies whether the MAX TNT functions as the user-side (terminal equipment) or network side (network terminating equipment) for T1 or E1 ISDN connections.

Usage: Specify one of the following values:

- `te` specifies the user side.
- `nt` specifies the network side.

Example: `set isdn-emulation-side = nt`

Dependencies: To specify `nt` for E1 connections, you must first set the Switch-Type parameter to `net5-pri`.

Location: T1/{ any-shelf any-slot 0 }:line-interface, E1/{ any-shelf any-slot 0 }:line-interface

ISDN NFAS support for Japanese switch types

To introduce non-facility associated signaling (NFAS) support for Japanese switches, TAOS now supports implicit identification of the primary D-channel interface and explicit identification of all other interfaces, as required by Japanese switches. In previous releases, the use of the Channel Identification information element by Japanese switches prevented NFAS configurations from working.

Following is an example of PRI/T1 line configuration for NFAS with a Japanese switch:

```
admin> read t1 { 1 1 1}
T1/{ shelf-1 slot-1 1} read

admin> set line-interface signaling-mode = isdn-nfas
admin> set line-interface switch-type = japan-pri
admin> set line-interface nfas-group-id = 0
admin> set line-interface nfas-id = 0
admin> set line-interface channel 24 channel = nfas-primary-d-channel
admin> write
T1/{ shelf-1 slot-1 1} written
```

Frame Relay PVCs over switched ISDN connections

With MAX TNT TAOS 8.0.0, the MAX TNT supports permanent virtual circuits (PVCs) over switched ISDN connections, referred to as *switched PVCs*. Switched PVCs are established in the same way as nailed PVCs: on the basis of an exchange of LMI frames and the occurrence of a number of events. However, instead of using nailed bandwidth, a switched PVC uses an

ISDN B channel that is brought up by an outgoing or incoming call. Switched PVCs can use channels on any slot card that works with the Hybrid Access (HDLC2 and HDLC2-EC) cards.

To establish a switched PVC by placing an outgoing call, the MAX TNT initiates the call in the usual way. When the call has been placed and the B channel is available, the system begins exchanging LMI frames to establish Frame Relay link operations, a process that can take several seconds. Once the link is up, it works just like a PVC with an access rate of 64 Kbps or 56 Kbps, depending on the ISDN configuration.

To establish a switched PVC by accepting an incoming call, CLID or DNIS authentication is required to enable the MAX TNT to begin using Frame Relay encapsulation before accepting the call. When the connection has been accepted, the MAX TNT follows the same procedure as described above for outgoing calls.

Overview of switched PVC settings

Following is the new parameter, shown with its default setting, relevant to switched PVCs:


```
[in FRAME-RELAY/spvc-1  
nailed-mode = off
```

The Nailed-Mode parameter in a Frame-Relay profile can now be set to `off`, indicating switched mode. When this parameter is set to `off`, the DLCI Connection profile must also specify a switched call type and a dial number, as well as a CLID or called number.

Examples of switched PVC configurations

Figure 18 shows PPP clients dialing into a MAX TNT to reach a customer premises equipment (CPE) router (10.9.8.7/24) that is accessible across Frame Relay.

Figure 18. Switched PVC to a Frame Relay switch



If both the Frame-Relay profile and the Connection profile for the DLCI interface specify a switched rather than nailed call, the MAX TNT brings up the interface as a switched connection on the basis of packet routing (as it usually does for a switched connection). If the Connection profile for the DLCI interface also specifies CLID or DNIS, the MAX TNT can also accept an incoming call from 10.9.8.7/24 to bring up the PVC.

Example of a Frame-Relay profile for a switched PVC

A Frame-Relay profile in switched mode can specify one of the same logical interfaces as a profile in nailed mode: NNI, UNI-DTE, and UNI-DCE. The details of these interface types are described in the *MAX TNT Network Configuration Guide*. The following commands configure a sample Frame-Relay profile for a switched connection to the switch in Figure 18:

```
admin> new frame-relay spvc-1
FRAME-RELAY/spvc-1 read

admin> set active = yes

admin> set nailed-mode = off

admin> set called-number-type = 2

admin> set switched-call-type = 64k-clear

admin> set link-mgmt = ansi-t1.617d

admin> list
[in FRAME-RELAY/spvc-1 (new) (changed)]
fr-name* = spvc-1
active = yes
nailed-up-group = 1
nailed-mode = off
called-number-type = 2
switched-call-type = 64k-clear
telephone-number = ""
billing-number = ""
transit-number = ""
link-mgmt = ansi-t1.617d
call-by-call-id = 0
link-type = dte
n391-val = 6
n392-val = 3
n393-val = 4
t391-val = 10
t392-val = 15
MRU = 1532
dceN392-val = 3
dceN393-val = 4
link-mgmt-dlci = dlci0

admin> write
FRAME-RELAY/spvc-1 written
```

Example of a Connection profile for a switched PVC

To enable the MAX TNT to place an outgoing call to establish the switched PVC, a Connection profile must specify a switched call type and a dial number. To enable the system to accept an incoming call to establish the switched PVC, the profile must specify a CLID or called number, and the MAX TNT must require CLID or DNIS. For details about CLID and DNIS authentication, see the *MAX TNT Network Configuration Guide*.

You can configure multiple DLCI interfaces over a single switched PVC by specifying the same Frame-Relay profile in each of the Connection profiles. The number of DLCI interfaces that can be supported on a switched ISDN call is limited only by system resources.

For example, the following commands configure the unit to require DNIS:

```
admin> read answer
ANSWER-DEFAULTS read

admin> set clid-auth-mode = dnis-require

admin> write
ANSWER-DEFAULTS written
```

The following commands configure a Connection profile to the CPE router shown in Figure 18, enabling both incoming and outgoing calls.

```
admin> new conn cpe-router
CONNECTION/cpe-router read

admin> set active = yes

admin> set encaps = frame-relay

admin> set dial-number = 853784

admin> set calledNumber = 3783

admin> set ip-options remote-address = 10.9.8.7/24

admin> set telco-options call-type = off

admin> set fr-options frame-relay-profile = spvc-1

admin> set fr-options dlci = 55


admin> write
CONNECTION/cpe-router written
```

Frame Relay switched virtual circuits (SVCs)

With MAX TNT TAOS 8.0.0, the MAX TNT supports Frame Relay switched virtual circuit (SVC) network configuration. An SVC is a point-to-point switched connection, which provides a lower cost, usage-based alternative to Frame Relay PVCs. SVCs provide an easier configuration for VCs throughout a Frame Relay network, and allow flexibility in rerouting VCs when equipment becomes unavailable. Like other types of switched connections, SVCs can be initiated by a dial-in or dial-out call.


A dial-in Frame Relay SVC terminates locally. The MAX TNT receives the call on a Frame Relay interface (a data link). An example of a terminating SVC is shown in Figure 30.

Figure 19. Terminating SVC on a Frame Relay interface



A dial-out SVC is initiated as an outbound call on a Frame Relay interface, on the basis of either an explicit dial-out or IP routing. Figure 20 shows a Pipeline unit dialing into the MAX TNT using PPP or some other type of encapsulation. The MAX TNT establishes the inbound call and then dials out on a Frame Relay interface on the basis of IP routing, just as it would for another type of switched dial-out call.

Figure 20. Dial-out SVC on a Frame Relay interface



Unlike permanent virtual circuits (PVCs), which are nailed connections, SVCs are on-demand connections and must use E.164 addresses (ISDN numbers) to identify and route to the SVC interface. For a dial-out SVC, the address is the Dial-Number in a Connection or RADIUS profile. For a dial-in SVC, the address can be specified in the Frame-Relay profile or as the CLID in a Connection or RADIUS profile. Dial-in SVCs are CLID authenticated.

To set up an SVC, you must configure SVC options in these locations:

- Frame-Relay profile, for the data link interface associated with a physical T1 or E1 port
- Connection profile, to establish the switched connection on the Frame Relay interface

Current limitations

In this release, the Frame Relay SVC implementation is subject to the following limitations:

- For SVCs, the MAX TNT operates as a Frame Relay user-side device (DTE). Network-side operations are not currently supported.
- The ability to request a specific DLCI value for an SVC is not implemented.

Hardware requirements

The following card types are supported for SVCs:

- Hybrid Access (used with channelized T1, T3, or E1 cards)
- Unchannelized T1 (T1 Frameline)
- Unchannelized E1 (E1 Frameline)
- Serial WAN (SWAN)

Overview of SVC settings in a Frame Relay profile

For the system to bring up an SVC connection, the data link interface must be operating and configured properly, with SVC (Q.933) signaling enabled. The system initiates the Q.933 signaling sequence when demand for an SVC occurs. All Q.933 call control information is transmitted over DLCI 0, which must also be used for the link management protocols if LMI is in use. (For SVCs as for PVCs, the LMI setting must match that of the far end switch. However, LMI is not required.)

For details about configuring Frame Relay data link interfaces, see the *MAX TNT Network Configuration Guide*. As for any Frame-Relay profile, you must specify a name and set the

active parameter to yes. In addition, the following parameters (shown with sample values) are relevant to SVC configurations:

```
[in FRAME-RELAY/" "]
nailed-up-group = 9
switched-call-type = 64k-clear
link-mgmt-dlci = dlci0

[in FRAME-RELAY/"":svc-options]
enabled = yes
fr-address = 5085551234
```

Parameter	Specifies
Nailed-Up-Group	Group number assigned to nailed channels in a line profile, such as a T1 or E1 profile. The default is 1. For channels on a nailed T1 connection, make sure that the number of channels the MAX TNT uses for the link matches the number of channels used by the device at the other end of the link. In addition, make sure that only one T1 profile specifies the number to be used by the Frame Relay data link.
Switched-Call-Type	Type of bearer channel capability. For a T1 line set for ESF signaling and B8ZS encoding, the remote switch or router typically requires that you set this parameter to 64k-Clear. A setting of 56k-clear (the default) is required if the line is set to D4 signaling and AMI encoding. E1 lines typically use 64k-Clear.
Link-Mgmt-DLCI	DLCI to use for LMI link management on the Frame Relay data link. Valid values are dlci0 (the default) and dlci1023. When SVC signaling is enabled, the data link can use either ANSI or CCITT LMI, but Link-Mgmt-DLCI <i>must</i> be set to its default dlci0 value.
Enabled	Enable/disable SVC signaling (Q.933) on the Frame Relay data link. SVC signaling is disabled by default. Note that a single data link interface can support both a PVC and SVC configuration.
FR-Address	E.164 address for this data link. This is the CLID for dial-out SVC connections on this interface. E.164 addresses are ISDN numbers, including telephone numbers. E.164 addresses can contain up to 15 ASCII digits. For example, standard 10-digit United States telephone numbers, such as 508-555-1234, are native E.164 addresses.

Overview of SVC settings in Connection profiles

You can configure multiple Connection profiles over a single SVC-enabled Frame Relay data link by specifying the same Frame-Relay profile in each of the profiles. If the FR-Address is specified in the Frame-Relay profile, the Connection profiles do not have to set the CLID parameter.

The Dial-Number of a Connection profile must be set in each Connection profile. For both outgoing and incoming circuit establishment requests, the Dial-Number field contains the E.164 address of the remote station. The combination of this field and the subaddress (if required) must be a unique value.

The local E.164 address is typically specified in the FR-Address parameter in a Frame-Relay profile. This address can be specified by the CLID setting in a Connection profile if it is

different from the FR-Address number. An E.164 address specified in the CLID setting of a Connection profile overrides the value of the FR-Address.

For details about configuring Frame-Relay encapsulated Connection profiles, see the *MAX TNT Network Configuration Guide*. As for any Connection profile, you must specify a station name and set the `active` parameter to `yes`. In addition, the following parameters (shown with sample values) are relevant to SVC configurations:

```
[in CONNECTION/svc-cx]
encapsulation-protocol = frame-relay
called-number-type = international
dial-number = 1235551212
clid = ""
subaddress = ""

[in CONNECTION/svc-cx:ip-options]
remote-address = 1.1.1.10/8

[in CONNECTION/svc-cx:telco-options]
data-service = frame-relay-svc

[in CONNECTION/svc-cx:fr-options]
frame-relay-profile = stdx-svc2
circuit-type = svc
dlci = 16
```


Parameter	Specifies
Encapsulation-Protocol	Encapsulation method for the connection. This parameter must be set to <code>frame-relay</code> for SVC connections.
Called-Number-Type	This field is set automatically to <code>international</code> when you write a Connection profile that has <code>circuit-type</code> set to <code>svc</code> .
Dial-Number	E.164 address of the remote station. E.164 addresses are ISDN numbers, including telephone numbers. E.164 addresses can contain up to 15 ASCII digits. For example, standard 10-digit United States telephone numbers, such as 123-555-1212, are E.164 addresses. The combination of this field and the subaddress must be a unique value.
CLID	E.164 address of the local end of the SVC. The local E.164 address is typically specified in the FR-Address parameter in a Frame-Relay profile. If an E.164 address is specified in the CLID parameter, it overrides the value of the FR-Address. There is no restriction on specifying the same CLID in multiple Connection profiles.
Subaddress	Subaddress portion of the E.164 address of the remote station, if a subaddress is required.
Remote-Address	IP address of the remote station.
Data-Service	This parameter is set automatically to <code>frame-relay-svc</code> when you write a Connection profile that has <code>circuit-type</code> set to <code>svc</code> .
Frame-Relay-Profile	Name of the Frame-Relay profile for the data link connection.
Circuit-Type	Type of virtual circuit Setting the value to <code>svc</code> determines that the circuit is established via Frame Relay SVC call signaling when data transfer is required.

Parameter	Specifies
DLCI	The system ignores this parameter for a Connection profile that has <code>circuit-type</code> set to <code>svc</code> . For an SVC, the DLCI value is assigned to the circuit by the network. The range of DLCI values for circuits is shared between PVCs and SVCs, and is managed between the network and user entities.

Examples of configuring Frame Relay SVCs

In the example SVC setup shown in Figure 21, the two MAX TNT units each contain channelized T1 and Hybrid Access cards. The switch is configured for SVC operation on the two T1 lines.

Figure 21. SVC between MAX TNT units with an intervening Frame Relay switch



For details about configuring the T1 lines for a Frame Relay data link interface, see the *MAX TNT Network Configuration Guide*.

Configuring the near-end MAX TNT for a Frame Relay SVC

The following commands on the unit labeled MAX TNT-01 (Figure 21) configure a Frame Relay data link interface on a T1 line that uses nailed-group 7:

```
admin> new frame-relay stdx-svc1
FRAME-RELAY/stdx-svc1 read
admin> set active = yes
admin> set nailed-up-group = 7
admin> set switched-call-type = 64k-clear
admin> set link-mgmt = ansi-t1.617d
admin> set svc-options enabled = yes
admin> set svc-options fr-address = 5085551234
admin> write -f
FRAME-RELAY/stdx-svc1 read
```

The following commands configure the SVC Connection profile to MAX TNT-02:

```
admin> new connection svc-555
CONNECTION/svc-555 read
admin> set active = yes
admin> set encapsulation-protocol = frame-relay
admin> set dial-number = 1235551212
```

```
admin> set ip-options remote-address = 3.3.3.3/24
admin> set fr-options frame-relay-profile = stdx-svc1
admin> set fr-options circuit-type = svc
admin> write -f
CONNECTION/svc-555 written
```

Configuring the far-end MAX TNT for a Frame Relay SVC

The following commands on the unit labeled MAX TNT-02 (Figure 21) configure a Frame Relay data link interface on a T1 line that uses nailed-group 8:

```
admin> new frame-relay stdx-svc2
FRAME-RELAY/stdx-svc2 read
admin> set active = yes
admin> set nailed-up-group = 8
admin> set switched-call-type = 64k-clear
admin> set link-mgmt = ansi-t1.617d
admin> set svc-options enabled = yes
admin> set svc-options fr-address = 1235551212
admin> write -f
FRAME-RELAY/stdx-svc2 read
```

The following commands configure the SVC Connection profile to MAX TNT-01:

```
admin> new connection svc-937
CONNECTION/svc-937 read
admin> set active = yes
admin> set encapsulation-protocol = frame-relay
admin> set dial-number = 5085551234
admin> set ip-options remote-address = 1.1.1.1/24
admin> set fr-options frame-relay-profile = stdx-svc2
admin> set fr-options circuit-type = svc
admin> write -f
CONNECTION/svc-937 written
```

Multilink Frame Relay (MFR)

Multilink Frame Relay enables the MAX TNT to bundle multiple Frame Relay data links to appear as a single logical data link with the aggregate bandwidth of the individual links. The bundled links are referred to as an *MFR bundle*. The bandwidth in an MFR bundle must be nailed, and can reside on a FrameLine (UT1), E1 FrameLine (UE1), or Hybrid Access card. All member data links of an MFR bundle must reside on the same card.


Note: MFR is supported on Hybrid Access cards, which can be used with T1, E1, or T3 cards to implement an MFR bundle. However, if more than one Hybrid Access card is installed in the system, the aggregate bandwidth of the bundle must be bound to a single Hybrid Access card by means of Call-Route profiles. If member data links of an MFR bundle span Hybrid Access cards, the link will fail. For details, see “Call routing requirements with multiple Hybrid Access cards” on page 189.

For background information about configuring Frame Relay in the MAX TNT, see the *MAX TNT Network Configuration Guide*.

DTE-DTE aggregation

Currently, end-to-end (DTE-DTE) aggregation is supported, which enables two DTEs to use the aggregate bandwidth of an MFR bundle across a regular Frame Relay (non-MFR) network. The fact that aggregate bandwidth of multiple links is in use is transparent to the Frame Relay switching equipment that resides between MFR peers. Figure 22 shows two DTEs using an MFR bundle of three data links through a Frame Relay network.

Figure 22. MFR DTE-DTE aggregation



To aggregate the bandwidth, the MAX TNT uses a segmentation-sequencing-reassembly protocol described in the Frame Relay Fragmentation Implementation Agreement FRF.12, which is based on the Multilink PPP (MP) protocol described in RFC 1990.

Current limitations

In this release, the MFR implementation is subject to the following limitations:

- The member data links of an MFR bundle must reside on the same slot card. This requirement is the only limitation on the number of links in a bundle.
- End-to-end fragmentation and reassembly are not supported.
- MFR using SVCs or switched PVCs is not supported.

How MFR bundles work


To create an MFR bundle, you define the Frame-Relay data link profiles as usual, but also specify an MFR bundle name (the name of a Multi-Link-FR profile) to make the data link a member of the specified bundle. The only limitation on the number of data links in a bundle is that the entire aggregate nailed bandwidth must all reside on the same slot card.

The member data links can provide different amounts of bandwidth. However, bundling data links that provide different amounts of bandwidth could result in throughput that is less than the sum of the member data links, because packets are sent to each of the members in a round-robin fashion without taking bandwidth into account. (For example, if an MFR bundle includes two data links on full T1 lines and one on fractional T1, some throughput might be lost due to packet queuing on the full T1 data links.)

Each data link within the bundle also requires at least one DLCI interface to the far-end device (the MFR peer). You must define the bundle first, before creating DLCI interfaces to the peer.

Figure 23 shows three bundled data links going through the Frame Relay network. Each data link has two DLCIs: 16 and 17. Data for each DLCI is sent to each of the member data links in a round-robin fashion.

Figure 23. MFR peers with three data links supporting two DLCIs



Because the DTE-DTE PVC goes through a non-MFR network, all of the individual links support the full User Network Interface (UNI) standards. As long as one DLCI from any of the bundled data links is active, that DLCI is considered active to the higher layers. For example, if data link 1 is down and DLCI 16 in data link 2 is active, the MFR peers (CPE-1 and CPE-2) consider DLCI 16 to be active.

Call routing requirements with multiple Hybrid Access cards

To implement an MFR bundle using a T1, E1, or T3 card with an Hybrid Access card, you must ensure that the aggregate bandwidth is bound to the channels of a single Hybrid Access card. If the system supports only one Hybrid Access card, no call routing configuration is required. However, if more than one Hybrid Access card is installed, you must define Call-Route profiles to map the bandwidth of the MFR bundle to the same Hybrid Access card.

The following sample output shows a system with E1 cards and Hybrid Access cards:

```
admin> show
Shelf 1 ( standalone ) :
  { shelf-1 slot-1 0 }      UP      8e1-card
  { shelf-1 slot-2 0 }      UP      8e1-card
  { shelf-1 slot-3 0 }      UP      hdlc2-card
  { shelf-1 slot-4 0 }      UP      hdlc2-card
  { shelf-1 slot-15 0 }     UP      8e1-card
  { shelf-1 slot-16 0 }     UP      8e1-card
```

Note: Because one Hybrid Access card can provide 186 channels (31 x 6) for MFR, one Hybrid Access card can support up to six Call-Route profiles that bind its channels to up to six back-to-back E1 ports. This setup places a limitation on the size of the MFR bundle when you are using an Hybrid Access card.

Example with two E1 lines in an MFR bundle

In the following example, the administrator creates two Call-Route profiles for the Hybrid Access card in slot 3, with each profile binding 31 HDLC channels to a single E1 line on the card in slot 2. The default Call-Route profile for the Hybrid Access card can be deleted or left unmodified, but must not be modified to specify an explicit route.

For example, the following commands create a Call-Route profile for the Hybrid Access card in slot 3 and set the preferred source to the first E1 interface in slot 2:

```

admin> new call-route { { { shelf-1 slot-3 0 } 0 } 1 }
CALL-ROUTE/{ { { shelf-1 slot-3 0 } 0 } 1 } read

admin> set preferred-source = { { 1 2 1 } 0 }

admin> list
[in CALL-ROUTE/{ { { shelf-1 slot-3 0 } 0 } 1 } (new) (changed)]
index* = { { { shelf-1 slot-3 0 } 0 } 1 }
trunk-group = 0
telephone-number = ""
preferred-source = { { shelf-1 slot-2 1 } 0 }
call-route-type = digital-call-type

admin> write
CALL-ROUTE/{ { { shelf-1 slot-3 0 } 0 } 1 } written

```

The next set of commands creates another Call-Route profile for the Hybrid Access card and sets the preferred source to the second E1 interface in slot 2:

```

admin> new call-route { { { shelf-1 slot-3 0 } 0 } 2 }
CALL-ROUTE/{ { { shelf-1 slot-3 0 } 0 } 1 } read

admin> set preferred-source = { { 1 2 2 } 0 }

admin> write
CALL-ROUTE/{ { { shelf-1 slot-3 0 } 0 } 2 } written

```

Note that the default Call-Route profile for the Hybrid Access card was not modified. It still specifies a general route for the card as a whole, as shown in the following listing:

```

admin> get call-route { { { shelf-1 slot-3 0 } 0 } 0 }
[in CALL-ROUTE/{ { { shelf-1 slot-3 0 } 0 } 0 } ]
index* = { { { shelf-1 slot-3 0 } 0 } 0 }
trunk-group = 0
telephone-number = ""
preferred-source = { { any-shelf any-slot 0 } 0 }
call-route-type = digital-call-type

```

Example with six E1 lines in an MFR bundle

If the MFR bundle aggregates enough bandwidth to utilize all of the channels on an Hybrid Access card (up to 186, or six E1 lines), you can create a single Call-Route profile that maps the E1 card to the Hybrid Access card. Only six of the E1 lines are usable for MFR, however.

For example, the following commands modify the default Call-Route profile to specify the E1 card in slot 2 as the preferred source for the card:

```

admin> read call-route { { { shelf-1 slot-3 0 } 0 } 0 }
CALL-ROUTE/{ { { shelf-1 slot-3 0 } 0 } 0 } read

admin> set preferred-source = { { 1 16 0 } 0 }

admin> list
[in CALL-ROUTE/{ { { shelf-1 slot-3 0 } 0 } 0 } (changed)]
index* = { { { shelf-1 slot-3 0 } 0 } 0 }
trunk-group = 0
telephone-number = ""
preferred-source = { { shelf-1 slot-16 0 } 0 }
call-route-type = digital-call-type

admin> write
CALL-ROUTE/{ { { shelf-1 slot-3 0 } 0 } 0 } written

```

Overview of MFR settings

Following are the parameters used for defining a bundle, shown with their default settings:

```
[in FRAME-RELAY/" "]
link-mgmt = none
link-type = dte
mfr-bundle-name = ""


[in MULTI-LINK-FR/" "]
mfr-bundle-name* = ""
active = no
mfr-bundle-type = mfr-dte
max-bundle-members = 1
min-bandwidth = 0
```

Parameter	Specifies
Link-Mgmt	Link management protocol. Settings are None (the default, which disables link management), ANSI-T1.617 (Annex D), and CCITT-Q.933a (CCITT Q.933 Annex A). A setting of none is not recommended for MFR. If the MAX TNT is connected to a Frame Relay switch, set the management protocol to the value used by the switch. If it is connected back-to-back to another MAX TNT unit, set it to the protocol used by the MFR peer.
Link-Type	Link-Type can be set to dte or dce for bundled data links. In this release, the nni setting is not supported for MFR.
MFR-Bundle-Name	Name of the bundle. The name can contain up to 15 characters and must be unique system-wide. In the Frame-Relay profile, the name binds the data link to an MFR bundle. All member data links must specify the same bundle name.
Active	Enable/disable the profile for use.
MFR-Bundle-Type	Type of MFR configuration. In this release, only the MFR-DTE configuration is supported. Future releases will support MFR UNI/NNI.
Max-Bundle-Members	Maximum number of data links allowed to join the MFR bundle. The default value is 1. All member data links must reside on the same slot card, so the card's capacity imposes a practical limitation on the maximum number of bundle members. If set to a number higher than 1, you can add bandwidth to the bundle up to the specified number of data links. For example, if Max-Bundle-Members is set to 4 and the bundle has 2 data links, you can add bandwidth dynamically by configuring another data link profile with the bundle name.
Min-Bandwidth	Minimum aggregated bandwidth before the bundle is considered inactive. In this release, you must leave the default zero value. Because of an unresolved problem in Frame Relay, if Min-Bandwidth is set to any other value, data is not sent on the bundle.

Example of an MFR DTE-DTE configuration

Figure 24 shows two MAX TNT units acting as MFR peers across the Frame Relay network. Each unit has two data links, each of which support two DLCI interfaces.

Figure 24. Sample MFR configuration



In each of the MFR peers, the bandwidth used by the bundled data links must reside on the same card. For each Frame-Relay data link profile in a bundle, you must also define a DLCI interface. Connection profiles for DLCI interfaces on bundled data links must specify the same remote IP address (that of the MFR peer), but must specify different DLCI numbers and Frame-Relay profiles.

Note: MFR configurations require multiple Connection profiles with the same Remote-Address. For most other types of configuration, the system prevents this condition. To save a Connection profile that specifies the same address as an existing Connection profile, make sure that the Frame-Relay profile it refers to already exists and specifies a bundle name.

Configuring MFR on CPE-1 using FrameLine

On CPE-1, the following commands create an MFR bundle consisting of two data links on a FrameLine (UT1) card:

```
admin> new frame-relay ut1.3-fr
FRAME-RELAY/ut1.3-fr read

admin> set active = yes

admin> set link-type = dte

admin> set nailed-up-group = 10

admin> set link-mgmt = ccitt

admin> set mfr-bundle-name = ut1-mfr

admin> write
FRAME-RELAY/ut1.3-fr written

admin> new frame-relay ut1.8-fr
FRAME-RELAY/ut1.8-fr read

admin> set active = yes

admin> set link-type = dte

admin> set nailed-up-group = 11

admin> set link-mgmt = ccitt

admin> set mfr-bundle-name = ut1-mfr

admin> write
FRAME-RELAY/ut1.8-fr written
```



```
admin> new multi-link-fr ut1-mfr
MULTI-LINK-FR/ut1-mfr read

admin> set active = yes

admin> set max-bundle-members = 2

admin> write
MULTI-LINK-FR/ut1-mfr written
```

The following commands on CPE-1 create DLCI interfaces on the bundled data links:

```
admin> new conn mfr1
CONNECTION/mfr1 read

admin> set active = yes

admin> set encaps = frame-relay

admin> set ip-options remote-address = 2.2.2.2/24

admin> set telco-options call-type = ft1

admin> set fr-options frame-relay-profile = ut1.3-fr

admin> set fr-options dlci = 100

admin> write
CONNECTION/mfr1 written

admin> new conn mfr2
CONNECTION/mfr2 read

admin> set active = yes

admin> set encaps = frame-relay

admin> set ip-options remote-address = 2.2.2.2/24

admin> set telco-options call-type = ft1

admin> set fr-options frame-relay-profile = ut1.8-fr

admin> set fr-options dlci = 200

admin> write
CONNECTION/mfr2 written
```

Configuring MFR on CPE-2 using T1

On CPE-2, the following commands create an MFR bundle of two data links that use lines 7 and 8 of a T1 card:

```
admin> new frame-relay ct1.7-fr
FRAME-RELAY/ct1.7-fr read

admin> set active = yes

admin> set link-type = dte

admin> set nailed-up-group = 10

admin> set link-mgmt = ccitt

admin> set mfr-bundle-name = ct1-mfr

admin> write
FRAME-RELAY/ct1.7-fr written

admin> new frame-relay ct1.8-fr
FRAME-RELAY/ct1.8-fr read

admin> set active = yes
```

```
admin> set link-type = dte
admin> set nailed-up-group = 11
admin> set link-mgmt = ccitt
admin> set mfr-bundle-name = sds1-mfr
admin> write
FRAME-RELAY/ctl1.8-fr written
admin> new multi-link-fr ctl1-mfr
MULTI-LINK-FR/ctl1-mfr read
admin> set active = yes
admin> set max-bundle-members = 2
admin> write
MULTI-LINK-FR/ctl1-mfr written
```

The following commands on CPE-2 specify DLCI interfaces on the bundled links:

```
admin> new conn mfr1
CONNECTION/mfr1 read
admin> set active = yes
admin> set encaps = frame-relay
admin> set ip-options remote-address = 1.1.1.1/24
admin> set telco-options call-type = ft1
admin> set fr-options frame-relay-profile = ctl1.7-fr
admin> set fr-options dlci = 100
admin> write
CONNECTION/mfr1 written
admin> new conn mfr2
CONNECTION/mfr2 read
admin> set active = yes
admin> set encaps = frame-relay
admin> set ip-options remote-address = 1.1.1.1/24
admin> set telco-options call-type = ft1
admin> set fr-options frame-relay-profile = ctl1.8-fr
admin> set fr-options dlci = 200
admin> write
CONNECTION/mfr2 written
```

Commands for monitoring MFR bundles

The FrameLine (UT1), E1 FrameLine (UE1), and Hybrid Access (HDLC2 and HDLC2-EC) cards on which MFR bandwidth resides support the following commands for monitoring MFR:

- Mfrdump command for displaying the bundle contents and the contents of its member data links
- Mfrdbg command for displaying MFR-related debug-level messages during data transmission on the MFR link

To use these commands, you must first open a session with the card that supports the nailed MFR bandwidth.

Mfrdump command

To dump the contents of an MFR bundle, open a session with the card and enter the Mfrdbg command. For example:

```
admin> open 1 4

ut1-1/4> mfrdump
ut1-mfr: next: 0 active: YES type: 0 maxMembers: 2
minBandwidth: 0 currentBandwidth: 0 activeCount: 2
nextSender: e01e6610 members: e01dee10 nextSeqIn: 3308 nextSeqOut:
3308
framesQueued: 0 buffer: 0

Member: 21 next: e01e6610 bandwidth: 0 lastSeq: -1
Member: 17 next: e01deb90 bandwidth: 0 lastSeq: -1
```

The command output includes the name of the MFR bundle (ut1-mfr in the sample output), and the following information about the bundle:

- The `next` field contains 0 if only one bundle is supported on the card. If the value is not zero, it indicates the address of the next MFR bundle on the card.
- The `type` field contains a number indicating the type of MFR aggregation. A 0 (zero) indicates DTE-DTE, which is the only value currently supported.
- The `maxMembers` field shows the maximum number of member data links.
- The bandwidth fields do not reflect accurate bandwidth measurements, but the minimum bandwidth must always be zero in this release.
- The `activeCount` value is the number of active member data links in the bundle. It can range from 1 to the value of `maxMembers`.
- The `nextSender` field points to the next member data link to send a packet out. The `members` field contains a pointer to the start of the member list. An unusual number in the pointer fields (for example, 1 or FFFFFFFF) can indicate a problem with the bundle.
- The `nextSeqIn` and `nextSeqOut` fields show the next sequence number expected (to be received and sent, respectively). Values range from 0 to 0xFFFF.
- The `framesQueued` value specifies the number of frames received from the Frame Relay link and waiting to be forwarded to the protocol level. Frames are queued only when they are received out of sequence. The `buffer` value is a pointer to the start of the frames queued buffer. A 0 means that no frames are queued.

For each of the member data links, the following additional information is provided:

- The `Member` field contains the DLCI number for this member.
- The `next` field contains a pointer to the next member data link in the list.
- The `bandwidth` field does not reflect accurate bandwidth measurements in this release.
- The `lastSeq` field shows the last sequence number queued on this member. The value will be -1 if no frames are queued.

Mfrdbg command

After beginning a multilink data transmission, open a session with the card, enable debug, and enter the Mfrdbg command to toggle display of the messages. For example:

```
admin> open 1 8

utl-1/8> debug on
Diagnostic output enabled

utl-1/8> mfrdbg
MFR debug is ON
```

The following output shows the system adding member data links to a bundle. It then shows the receive/send responses that indicate data is being transmitted over the link.

```
utl-1/8> mfrAddActiveMember: bundleName is utl-mfr
_mfr Find Bundle: bundle not found.
mfrAddActiveMember: Bundle utl-mfr not found!
_mfr Create Bundle: bundle name is utl-mfr
mfrAddActiveMember: Adding Member
mfrAddActiveMember: bundleName is utl-mfr
mfrAddActiveMember: Adding Member
mfrAddActiveMember: bundleName is utl-mfr
mfrAddActiveMember: Adding Member
mfrAddActiveMember: bundleName is utl-mfr
mfrAddActiveMember: Adding Member
mfrAddActiveMember: bundleName is utl-mfr
mfrAddActiveMember: Adding Member
mfrAddActiveMember: bundleName is utl-mfr
mfrAddActiveMember: Adding Member
mfrDataRecieve: header is: 0000c100
mfrDataRecieve: sequenceNumber: 0
mfrDataRecieve: Pass Data on
mfrDataReceive: Sequence: 0
MFR header is: : 2 of 88 octets
E024898C: c1 00 ..
mfrDataRecieve: header is: 0000c101
mfrDataRecieve: sequenceNumber: 1
mfrDataRecieve: Pass Data on
mfrDataReceive: Sequence: 1
MFR header is: : 2 of 88 octets
E024898C: c1 01 ..
mfrDataRecieve: header is: 0000c102
mfrDataRecieve: sequenceNumber: 2
mfrDataRecieve: Pass Data on
mfrDataReceive: Sequence: 2
MFR header is: : 2 of 88 octets
E024898C: c1 02 ..
mfrDataRecieve: header is: 0000c103
mfrDataRecieve: sequenceNumber: 3
mfrDataRecieve: Pass Data on
mfrDataReceive: Sequence: 3
MFR header is: : 2 of 88 octets
E024898C: c1 03 ..
mfrDataRecieve: header is: 0000c104
mfrDataRecieve: sequenceNumber: 4
mfrDataRecieve: Pass Data on
mfrDataReceive: Sequence: 4
MFR header is: : 2 of 88 octets
E024898C: c1 04 ..
mfrDataRecieve: header is: 0000c105
mfrDataRecieve: sequenceNumber: 5
mfrDataRecieve: Pass Data on
mfrDataReceive: Sequence: 5
MFR header is: : 2 of 88 octets
```

```
E024898C: c1 05 ..  
mfrDataRecieve: header is: 0000c106  
mfrDataRecieve: sequenceNumber: 6  
mfrDataRecieve: Pass Data on  
mfrDataReceive: Sequence: 6  
MFR header is: : 2 of 88 octets  
E024898C: c1 06 ..
```

Parameter reference entries

MFR-Bundle-Name

Description: Specifies the name of a multilink Frame Relay bundle. The name can contain up to 15 characters and must be unique system-wide.

In a Multi-Link-FR profile, this parameter defines a name for the bundle. In a Frame-Relay profile, it makes the data link a member of the MFR bundle. All member data links must specify the same bundle name in the Frame-Relay profile.

Usage: Specify the name of a Multi-Link-FR profile (up to 15 characters).

Example: `set mfr-bundle-name = sds1-mfr`

Dependencies: All member data links must specify the name of the same Multi-Link-FR profile.

Location: Frame-Relay, Multi-Link-FR

See Also: Active, MFR-Bundle-Type, Max-Bundle-Members, Min-Bandwidth

Active

Description: Enable/disable the profile for use.

Usage: Specify yes or no. The default is no.

Example: `set active = yes`

Location: Multi-Link-FR

See Also: MFR-Bundle-Name, MFR-Bundle-Type, Max-Bundle-Members, Min-Bandwidth

MFR-Bundle-Type

Description: Specifies the type of MFR configuration. In this release, only the MFR-DTE type is supported.

Usage: Specify `mfr-dte`.

Example: `set mfr-bundle-type = mfr-dte`

Location: Multi-Link-FR

See Also: MFR-Bundle-Name, Active, Max-Bundle-Members, Min-Bandwidth

Max-Bundle-Members

Description: Specifies the maximum number of data links allowed to join the MFR bundle.

Usage: Specify the maximum number of data links allowed to join the MFR value. The default value is 1. If set to a number higher than 1, you can add bandwidth to the bundle up to the specified number of data links. For example, if Max-Bundle-Members is set to 4 and the bundle has 2 data links, you can add bandwidth dynamically by configuring another data link profile with the bundle name.

Example: `set max-bundle-members = 4`

Dependencies: All member data links must reside on the same slot card, so the card's capacity imposes a practical limitation on the maximum number of bundle members.

Location: Multi-Link-FR

See Also: MFR-Bundle-Name, MFR-Bundle-Type, Active, Min-Bandwidth

Min-Bandwidth

Description: Specifies the minimum aggregated bandwidth before the bundle is considered inactive.

Usage: In this release, you must leave the default zero value. Because of an unresolved problem in Frame Relay, if Min-Bandwidth is set to any other value, data is not sent on the bundle.

Example: `set min-bandwidth = 0`

Location: Multi-Link-FR

See Also: MFR-Bundle-Name, MFR-Bundle-Type, Active, Max-Bundle-Members

DLCI bundling in MFR

In the initial implementation of MFR (described in "Multilink Frame Relay (MFR)" on page 187), you can bundle multiple Frame Relay data links to appear as a single logical data link with the aggregate bandwidth of the individual links. This functionality is still supported. This additional MFR feature enables you to bundle DLCIs. DLCI bundling enables a more flexible use of physical lines, because a single line can support both bundled and nonbundled connections.

Related parameter settings

Following are related parameters that were introduced in the initial release of MFR, shown with sample values related to DLCI bundling:

```
[in MULTI-LINK-FR/bundle1]
mfr-bundle-name* = bundle1
max-bundle-members = 4

[in FRAME-RELAY/ut1.3-fr]
mfr-bundle-name = ""
```

The following parameter (shown with a sample value) is new for this feature, and enables DLCI bundling:

```
[in CONNECTION/dp1:fr-options]
mfr-bundle-name = bundle1
```

Parameter	Specifies
MFR-Bundle-Name	<p>Name of the bundle. The name can contain up to 15 characters and must be unique system-wide. In a Connection profile, specifying a bundle name adds the DLCI to an MFR bundle. In a Frame-Relay profile, specifying a name adds the data link itself and all DLCIs that make use of it to an MFR bundle.</p> <p>To enable a line to support both MFR and non-MFR links, the bundle name in a Frame-Relay profile must be null.</p> <p>All member links in a bundle must specify the same bundle name.</p>
Max-Bundle-Members	<p>Maximum number of Frame Relay connections allowed to join the MFR bundle. The default value is 1. If this parameter is set to a number higher than 1, you can add data links or DLCIs to the bundle up to the specified number of connections. For example, if Max-Bundle-Members is set to 4 and the bundle has 2 DLCIs, you can add two additional DLCIs dynamically by configuring Connection profiles with the same bundle name.</p>


The system checks first for a bundle name in a Connection profile. It checks for a bundle name in the Frame-Relay profile only if no bundle name is found in the Connection profile.

Example of bundling DLCIs by means of MFR

In the following example, a company has three offices connected via Frame Relay. Office A receives a lot of traffic from Office B, and Office C receives very little traffic from Office B. Office B has two T1 lines to Office A.

As shown in Figure 25, instead of installing a third fractional T1 line from Office B to Office C, you can include traffic destined for Office C on one of the existing T1 lines and define a bundle of two DLCIs to Office A.

Figure 25. Example of MFR on per-DLCI basis



Because very little traffic is sent to Office C, most of the bandwidth of the second T1 line is available for traffic to Office A. The examples that follow show how to configure the MFR peers in Office B and Office A.

In each of the MFR peers, the bandwidth used by the bundled connections must reside on the same card. All Connection profiles for bundled DLCI interfaces must specify the remote IP

address of the same MFR peer, but must specify different DLCI numbers and Frame-Relay profiles.

Note: MFR configurations require multiple Connection profiles with the same Remote-Address. For most other types of configuration, the system prevents this condition. To save a Connection profile that specifies the same address as an existing profile, make sure that either the other Connection profile or the associated Frame-Relay profile specifies a bundle name.

Sample configurations in first MFR peer

The following commands on the CPE in Office B define an MFR bundle and two Frame-Relay data link profiles. (The Frame-Relay profiles use the third and fourth T1 lines of the same UT1 card.)

```
admin> new multi-link-fr ut1-mfr
MULTI-LINK-FR/ut1-mfr read

admin> set active = yes

admin> set max-bundle-members = 2

admin> write
MULTI-LINK-FR/ut1-mfr written

admin> new frame-relay ut1.3-fr
FRAME-RELAY/ut1.3-fr read

admin> set active = yes

admin> set link-type = dte

admin> set nailed-up-group = 10

admin> set link-mgmt = ccitt

admin> write
FRAME-RELAY/ut1.3-fr written

admin> new frame-relay ut1.4-fr
FRAME-RELAY/ut1.4-fr read

admin> set active = yes

admin> set link-type = dte

admin> set nailed-up-group = 11

admin> set link-mgmt = ccitt

admin> write
FRAME-RELAY/ut1.4-fr written
```

The following commands on the CPE in Office B create two DLCI interfaces to Office A and add them to the MFR bundle:

```
admin> new conn a-1
CONNECTION/a-1 read

admin> set active = yes

admin> set encaps = frame-relay

admin> set ip-options remote-address = 2.2.2.2/24

admin> set telco-options call-type = ft1

admin> set fr-options frame-relay-profile = ut1.3-fr

admin> set fr-options dlci = 100
```



```
admin> set fr-options mfr-bundle-name = ut1-mfr
admin> write
CONNECTION/a-1 written
admin> new conn a-2
CONNECTION/a-2 read
admin> set active = yes
admin> set encaps = frame-relay
admin> set ip-options remote-address = 2.2.2.2/24
admin> set telco-options call-type = ft1
admin> set fr-options frame-relay-profile = ut1.4-fr
admin> set fr-options dlci = 200
admin> set fr-options mfr-bundle-name = ut1-mfr
admin> write
CONNECTION/a-2 written
```

The following commands on the CPE in Office B create a DLCI interface to Office C:

```
admin> new conn c-1
CONNECTION/c-1 read
admin> set active = yes
admin> set encaps = frame-relay
admin> set ip-options remote-address = 3.3.3.3/24
admin> set telco-options call-type = ft1
admin> set fr-options frame-relay-profile = ut1.3-fr
admin> set fr-options dlci = 300
admin> write
CONNECTION/c-1 written
```

The profile to Office C is not part of the bundle.

Sample configurations in second MFR peer

The following commands on the CPE in Office A define an MFR bundle and two Frame-Relay data link profiles. (The Frame-Relay profiles use the first and second T1 lines on the same T1 card.)

```
admin> new multi-link-fr t1-mfr
MULTI-LINK-FR/t1-mfr read
admin> set active = yes
admin> set max-bundle-members = 2
admin> write
MULTI-LINK-FR/t1-mfr written
admin> new frame-relay t1.1-fr
FRAME-RELAY/t1.1-fr read
admin> set active = yes
admin> set link-type = dte
admin> set nailed-up-group = 50
```

```
admin> set link-mgmt = ccitt
admin> write
FRAME-RELAY/t1.1-fr written
admin> new frame-relay t1.2-fr
FRAME-RELAY/t1.2-fr read
admin> set active = yes
admin> set link-type = dte
admin> set nailed-up-group = 60
admin> set link-mgmt = ccitt
admin> write
FRAME-RELAY/t1.2-fr written
```

The next commands on the CPE in Office A create two DLCI interfaces to Office B and add them to the MFR bundle:

```
admin> new conn b-1
CONNECTION/b-1 read
admin> set active = yes
admin> set encaps = frame-relay
admin> set ip-options remote-address = 1.1.1.1/24
admin> set telco-options call-type = ft1
admin> set fr-options frame-relay-profile = t1.1-fr
admin> set fr-options dlci = 100
admin> set fr-options mfr-bundle-name = t1-mfr
admin> write
CONNECTION/b-1 written
admin> new conn b-2
CONNECTION/b-2 read
admin> set active = yes
admin> set encaps = frame-relay
admin> set ip-options remote-address = 1.1.1.1/24
admin> set telco-options call-type = ft1
admin> set fr-options frame-relay-profile = t1.2-fr
admin> set fr-options dlci = 200
admin> set fr-options mfr-bundle-name = t1-mfr
admin> write
CONNECTION/b-2 written
```

The next set of commands on the CPE in Office A create a Connection profile to receive the incoming traffic from Office B destined for Office C:

```
admin> new conn c-1
CONNECTION/c-1 read
admin> set active = yes
admin> set encaps = frame-relay
admin> set ip-options remote-address = 1.1.1.1/24
```

```
admin> set telco-options call-type = ft1
admin> set fr-options frame-relay-profile = t1.1-fr
admin> set fr-options dlci = 300
admin> write
CONNECTION/c-1 written
```

When the MAX TNT receives IP packets on this connection, it passes them up to layer 3. The next commands specify a static route to Office C and create a Connection profile to the intervening Frame Relay CPE:

```
admin> new ip-route frcpe
IP-ROUTE/frcpe read
admin> set dest = 3.3.3.3/24
admin> set gateway = 5.5.5.5
admin> write
IP-ROUTE/frcpe written
admin> new conn frswitch
CONNECTION/frswitch read
admin> set active = yes
admin> set encapsulation-protocol = frame-relay
admin> set ip-options remote-address = 5.5.5.5/24
admin> set telco-options call-type = ft1
admin> set fr-options frame-relay-profile = sds1.7-fr
admin> set fr-options dlci = 400
admin> write
CONNECTION/frswitch written
```

Parameter reference entry

The meaning of the following two parameters has been modified slightly for DLCI bundling.

MFR-Bundle-Name

Name of a Multi-Link-FR profile that defines a bundle. The name can contain up to 15 characters, and must be unique system-wide. In a Frame-Relay profile, specifying a bundle name adds the data link itself and all DLCIs that make use of it to an MFR bundle. Specifying a bundle name in a Connection profile adds the DLCI to an MFR bundle.

To enable a line to support both MFR and non-MFR links, the bundle name in a Frame-Relay profile must be null.

Usage: Specify the name of a Multi-Link-FR profile (up to 15 characters).

Example: `set mfr-bundle-name = sds1-mfr`

Dependencies: All member links in a bundle must specify the same bundle name.

Location: Frame-Relay, Multi-Link-FR, Connection

See Also: Active, MFR-Bundle-Type, Max-Bundle-Members, Min-Bandwidth

Max-Bundle-Members

Description: Maximum number of Frame Relay connections allowed to join the MFR bundle.

Usage: Specify the maximum number of data links or DLCIs allowed to join the MFR bundle. The default value is 1. If set to a number higher than 1, you can add data links or DLCIs to the bundle up to the specified number of connections. For example, if Max-Bundle-Members is set to 4 and the bundle has 2 DLCIs, you can add two additional DLCIs dynamically by configuring Connection profiles with the same bundle name.

Example: `set max-bundle-members = 4`

Location: Multi-Link-FR

See Also: MFR-Bundle-Name, MFR-Bundle-Type, Active, Min-Bandwidth

MFR circuit switching

With MAX TNT TAOS 8.0.0, the Multilink Frame Relay (MFR) protocol supports Frame Relay switching by means of a circuit configuration. A Frame Relay circuit is a PVC segment configured in two Connection profiles. Data coming in on the DLCI configured in one Connection profile is switched to the DLCI configured in the paired profile.


MFR circuit switching allows you to configure a circuit that switches an MFR bundle to another MFR bundle, or to a single data link. Data coming in on an MFR bundle (multiple DLCIs) is switched to the other circuit endpoint, which can also support either a single DLCI or an MFR bundle.

Note: For MFR circuit switching, both sides of the circuit must be DLCI interfaces. ATM-Frame Relay circuits are not supported with MFR.

No new parameters or RADIUS attributes are required for MFR circuit switching. MFR bundles are configured as described in “Multilink Frame Relay (MFR)” on page 187. Circuits are configured as they were in previous releases, as described in the *MAX TNT Network Configuration Guide*.

For example, Figure 26 shows a MAX TNT unit that switches between an MFR bundle on one side and single data link interface on the other.

Figure 26. MFR circuit switching from a bundle to a single T1 interface



The next sections provide sample profiles on the unit labeled ABC-unit for the switching configuration shown in Figure 26.

Configuring the MFR bundle

The following commands define the MFR bundle that aggregates the bandwidth of the two T1 interfaces to XYZ-unit:

```
admin> new multi-link-fr ut1-mfr
MULTI-LINK-FR/ut1-mfr read

admin> set active = yes

admin> set max-bundle-members = 2

admin> write
MULTI-LINK-FR/ut1-mfr written
```

The following commands configure the two T1 data link interfaces to XYZ-unit:

```
admin> new frame-relay ut1.3-xyz
FRAME-RELAY/ut1.3-xyz read

admin> set active = yes

admin> set link-type = dte

admin> set nailed-up-group = 10

admin> set link-mgmt = ccitt

admin> set mfr-bundle-name = ut1-mfr

admin> write
FRAME-RELAY/ut1.3-xyz written

admin> new frame-relay ut1.8-xyz
FRAME-RELAY/ut1.8-xyz read

admin> set active = yes

admin> set link-type = dte

admin> set nailed-up-group = 11

admin> set link-mgmt = ccitt

admin> set mfr-bundle-name = ut1-mfr

admin> write
FRAME-RELAY/ut1.8-xyz written
```

Both of the sample Frame-Relay profiles shown immediately above specify the MFR bundle name. As described in “DLCI bundling in MFR” on page 198, the bundle name can be specified either in a Frame-Relay profile or a Connection profile.

Configuring the MFR circuit endpoint

Each DLCI interface requires its own Connection profile, so the MFR circuit endpoint in this example uses two profiles. Each of the profiles references the Frame-Relay profile of a member data link.

Note: Although DLCI interfaces that specify different data links can use the same DLCIs, the use of unique DLCIs for these interfaces makes troubleshooting easier. Using unique DLCIs for interfaces that specify different data links in an MFR bundle is recommended, but not required.

In this example, the MFR bundle name is specified in the Frame-Relay profiles, so it is not repeated in the Connection profiles. Each DLCI interface can specify the bundle name either in the associated Frame-Relay profile or in the Connection profile.

The following commands create the DLCI interfaces for the MFR circuit endpoint:

```
admin> read conn xyz-1
CONNECTION/xyz-1 read
admin> set active = yes
admin> set encaps = frame-relay-circuit
admin> set ip-options ip-routing-enabled = no
admin> set telco-options call-type = ft1
admin> set fr-options frame-relay-profile = ut1.3-xyz
admin> set fr-options dlci = 116
admin> set fr-options circuit-name = circuit1
admin> write
CONNECTION/xyz-1 written
admin> read conn xyz-2
CONNECTION/xyz-2 read
admin> set active = yes
admin> set encaps = frame-relay-circuit
admin> set ip-options ip-routing-enabled = no
admin> set telco-options call-type = ft1
admin> set fr-options frame-relay-profile = ut1.8-xyz
admin> set fr-options dlci = 117
admin> set fr-options circuit-name = circuit1
admin> write
CONNECTION/xyz-2 written
```

Configuring the Frame Relay side of the circuit

The following commands configure the T1 data link interface to the MAX-unit in Figure 26:

```
admin> new frame-relay t1-max
FRAME-RELAY/t1-max read
admin> set active = yes
admin> set link-type = dce
admin> set nailed-up-group = 22
admin> set link-mgmt = ccitt
admin> write
FRAME-RELAY/t1-max written
```

The following commands configure the Frame Relay circuit endpoint to the MAX-unit:


```
admin> read conn max-1
CONNECTION/max-1 read
admin> set active = yes
admin> set encaps = frame-relay-circuit
```

```
admin> set ip-options ip-routing-enabled = no
admin> set telco-options call-type = ft1
admin> set fr-options frame-relay-profile = t1-max
admin> set fr-options dlci = 200
admin> set fr-options circuit-name = circuit1
admin> write
CONNECTION/max-1 written
```

Configurable VPI-VCI ranges

The Segment Assembly/Reassembly (SAR) unit on the OC3-ATM and DS3-ATM cards supports a 16-bit virtual path identifier and virtual channel identifier (VPI-VCI) range. In earlier software versions, the 16-bit VPI-VCI pair was hard-coded with 4 bits for the VPI and 12 for the VCI, as shown in Figure 27:

Figure 27. 16-bit VPI-VCI range



Now, you can select the best combination of VPI and VCI bit sizes to fit the list of supported VPI-VCI pairs obtained from the network provider. The new values take effect as soon as you write the OC3-ATM or DS3-ATM profile. Following are the relevant parameters, shown with their default values:

```
[in DS3-ATM/{ shelf-1 slot-3 1 }:line-config]
vpi-vci-range = 0-15/32-4095
[in OC3-ATM/{ shelf-1 slot-1 1 }:line-config]
vpi-vci-range = 0-15/32-4095
```

Parameter	Specifies																											
VPI-VCI-Range	<p>Range of values in the virtual path identifier and virtual channel identifier (VPI-VCI) pair. The default setting of 0-15/32-4095 is the range of values that can be represented in a 4-bit VPI and 12-bit VCI. This setting is compatible with earlier releases. Following are the possible ranges and their relevant bit sizes:</p> <table><tr><th>Range</th><th># Of VPI bits</th><th># Of VCI bits</th></tr><tr><td>0-1/32-32767</td><td>1</td><td>15</td></tr><tr><td>0-3/32-16383</td><td>2</td><td>14</td></tr><tr><td>0-7/32-8191</td><td>3</td><td>13</td></tr><tr><td>0-15/32-4095</td><td>4</td><td>12</td></tr><tr><td>0-31/32-2047</td><td>5</td><td>11</td></tr><tr><td>0-63/32-1023</td><td>6</td><td>10</td></tr><tr><td>0-127/32-511</td><td>7</td><td>9</td></tr><tr><td>0-255/32-255</td><td>8</td><td>8</td></tr></table>	Range	# Of VPI bits	# Of VCI bits	0-1/32-32767	1	15	0-3/32-16383	2	14	0-7/32-8191	3	13	0-15/32-4095	4	12	0-31/32-2047	5	11	0-63/32-1023	6	10	0-127/32-511	7	9	0-255/32-255	8	8
Range	# Of VPI bits	# Of VCI bits																										
0-1/32-32767	1	15																										
0-3/32-16383	2	14																										
0-7/32-8191	3	13																										
0-15/32-4095	4	12																										
0-31/32-2047	5	11																										
0-63/32-1023	6	10																										
0-127/32-511	7	9																										
0-255/32-255	8	8																										

OAM loopback for DS3 ATM PVC fault management

With MAX TNT TAOS 8.0.0, MAX TNT units can detect the failure of an ATM permanent virtual circuit (PVC) on a DS3-ATM interface by using Operation, Administration, and Maintenance (OAM) F5 loopback. When it detects failure, the system clears the PVC, puts the interface in an inactive state, and attempts to reestablish the nailed connection.

If you specify an interface that has been configured for VC fault management as an argument to the `oamloopback` command, the command rejects the request with an error message.

Overview of Connection profile settings

Following are the relevant parameters, shown with default values:


```
[in CONNECTION/"":atm-options]
vc-fault-management = none
vc-max-loopback-cell-loss = 1
```

Parameter	Specifies
VC-Fault-Management	VC fault management type. When the parameter is set to <code>none</code> (the default), no fault management is performed on the VC. If the parameter is set to <code>segment-loopback</code> , the system sends an OAM F5 segment loopback cell to the remote device every 5 seconds. If the parameter is set to <code>end-to-end-loopback</code> , the system sends an OAM F5 end-to-end loopback cell to the remote device every 5 seconds.
VC-Max-Loopback-Cell-Loss	Number of consecutive loopback cells that can be lost before the system clears the connection. When a PVC is cleared, the interface is in an inactive state until the system can reestablish the connection. The default is 1.

Example of configuring an ATM VC with fault management

An example MAX TNT connection to a remote ATM switch on a DS3-ATM interface is shown in In Figure 28.

Figure 28. ATM permanent virtual circuit



The following commands configure the DS3-ATM interface:

```
admin> read atm-ds3 {1 2 1}
DS3-ATM/{ 1 2 1 } read
admin> set name = atm-sf
admin> set enabled = yes
```



```
admin> set line nailed-group = 101
admin> write
ATM-DS3/{ shelf-1 slot-2 1 } written
```

The following commands configure the ATM PVC with end-to-end loopback fault management:

```
admin> new connection atmswitch
CONNECTION/atmswitch read
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 2.2.2.2
admin> set telco-options call-type = ft1
admin> set telco-options nailed-up-group = 101
admin> set atm-options vpi = 12
admin> set atm-options vci = 42
admin> set atm-options vc-fault-management = end-to-end-loopback
admin> set atm-options vc-max-loopback-cell-loss = 5
admin> write
CONNECTION/atmswitch written
```

With these settings, the system establishes the nailed connection and sends an OAM F5 end-to-end loopback cell every 5 seconds. If it does not receive the loopback cell back for 5 consecutive intervals (25 seconds), it clears the PVC, puts the interface in an inactive state, and begin attempts to reestablish the nailed connection.

ATM-direct

With MAX TNT TAOS 8.0.0, MAX TNT units support ATM-direct for concentrating incoming PPP calls onto an ATM interface. ATM-direct aggregates multiple PPP connections and forwards them as a combined data stream solely on the basis of the ATM-Direct configuration. An upstream device then examines the packets and routes them appropriately.

Note: An ATM-direct connection is not a full-duplex tunnel between a PPP dial-in user and a far-end device. Although the MAX TNT does not route the packets onto the ATM link, it must use the router to send packets received across ATM back to the appropriate PPP caller. For this reason, ATM-direct connections must enable IP routing.

When a PPP Connection profile specifies ATM-Direct, the MAX TNT forwards the data stream out on a specified ATM virtual circuit (VC). It leaves the task of routing the packets to an upstream device.

Settings in a Connection profile

Following are the relevant ATM-direct parameters, shown with default settings:

```
[in CONNECTION/" "]
encapsulation-protocol = mpp

[in CONNECTION/" ":atm-options]
atm-direct-enabled = no
atm-direct-profile = "
```

```
[in CONNECTION/"":ip-options]
ip-routing-enabled = yes
remote-address = 0.0.0.0/0
address-pool = 0
```

Parameter	Specifies
Encapsulation-Protocol	The encapsulation protocol. Must be set to ppp, mp, or mpp for ATM-direct connections.
ATM-Direct-Enabled	Enable/disable ATM-direct mode for this connection.
ATM-Direct-Profile	Name of a Connection or RADIUS profile that specifies an ATM link with a VPI-VCI pair.
IP-Routing-Enabled	Enable/disable IP routing for this connection. Must be enabled for the MAX TNT to send data back to the appropriate PPP caller.
Remote-Address	PPP caller's IP address. As the MAX TNT receives return packets for many ATM-direct connections across the same ATM link, it uses this address to determine the PPP caller that receives the return packets.
Address-Pool	Number of the address pool from which to acquire an address. If the Remote-Address is null and pools have been configured, the system assigns an IP address dynamically. For details about configuring and using dynamic IP addresses, see the <i>MAX TNT Network Configuration Guide</i> .

Settings in a RADIUS profile


RADIUS uses the following attribute-value pairs for ATM-direct connections:

RADIUS attribute	Value
Framed-Protocol (7)	The encapsulation protocol. Must be set to PPP (1), MP (262), or MPP (256) for ATM-direct connections.
Ascend-ATM-Direct (76)	Enable/disable ATM-direct mode for this connection. ATM-Direct-No (0) is the default. Set to ATM-Direct-Yes (1) for ATM-direct connections.
Ascend-ATM-Direct-Profile (77)	Name of a profile that specifies an ATM link with a VPI-VCI pair.
Ascend-Route-IP (228)	Enable/disable IP routing for this connection. (IP is enabled by default.) If this attribute is present, it must be set to Route-IP-Yes to enable the MAX TNT to send data back to the appropriate PPP caller.
Framed-IP-Address (8)	PPP caller's IP address. As the MAX TNT receives return packets for many ATM-direct connections across the same ATM link, it uses this address to determine the PPP caller that receives the return packets. If the Framed-IP-Address attribute-value pair is missing from the RADIUS profile and pools have been configured, the system assigns an IP address dynamically. For details about configuring and using dynamic IP addresses, see the <i>MAX TNT Network Configuration Guide</i> .
Framed-IP-Netmask (9)	A subnet mask for Framed-IP-Address.

Examples of ATM-direct connections

In Figure 29, the MAX TNT forwards the data stream from two PPP dial-in hosts across the same ATM link:

Figure 29. ATM-direct concentrating PPP calls to an ATM interface



For details about configuring the OC3-ATM or DS3-ATM line interface, see the slot card configuration guides at <http://www.ascend.com/doclibrary>. After you register, you can view or download the guide.

The following set of commands configures the ATM link with a VPI-VCI pair.

```
admin> new connection atm-switch-1
CONNECTION/atm-switch-1 read
admin> set active = yes
admin> set encapsulation-protocol = atm
admin> set ip-options remote-address = 1.1.1.1
admin> set telco-options call-type = ft1
admin> set telco-options nailed-up-group = 99
admin> set atm-options vpi = 15
admin> set atm-options vci = 35
admin> write
CONNECTION/atm-switch-1 written
```

The following set of commands configures ATM-direct Connection profiles for the incoming calls. The name of the profile for the connection to the ATM switch in this example is atm-switch-1:

```
admin> new conn caller-1
CONNECTION/caller-1 read
admin> set active = yes
admin> set encapsulation-protocol = ppp
admin> set ppp-options recv-password = caller1*3
admin> set ip-options remote-address = 10.5.6.7/32
admin> set atm-options atm-direct-enabled = yes
admin> set atm-options atm-direct-profile = atm-switch-1
admin> write
CONNECTION/caller-1 written
admin> new conn caller-2
CONNECTION/caller-2 read
```

```
admin> set active = yes
admin> set encapsulation-protocol = ppp
admin> set ppp-options recv-password = caller2!!8
admin> set ip-options remote-address = 10.7.8.9/32
admin> set atm-options atm-direct-enabled = yes
admin> set atm-options atm-direct-profile = atm-switch-1
admin> write
CONNECTION/caller-2 written
```

Following are comparable RADIUS profiles:

```
caller-1 Password = "caller1*3", Service-Type = Framed-User
      Framed-Protocol = PPP,
      Framed-IP-Address = 10.5.6.7,
      Framed-IP-Netmask = 255.255.255.255,
      Ascend-ATM-Direct = ATM-Direct-Yes,
      Ascend-ATM-Direct-Profile = "atm-switch-1"

caller-2 Password = "caller2!!8", Service-Type = Framed-User
      Framed-Protocol = PPP,
      Framed-IP-Address = 10.7.8.9,
      Framed-IP-Netmask = 255.255.255.255,
      Ascend-ATM-Direct = ATM-Direct-Yes,
      Ascend-ATM-Direct-Profile = "atm-switch-1"
```

Parameter reference entries

ATM-Direct-Enabled

Description: Specifies whether ATM-direct is enabled.

Usage: Specify yes or no. The default is no.

- yes specifies that ATM-direct is enabled.
- no specifies that ATM-direct is disabled.

Example: `set atm-direct-enabled = yes`

Location: Connection > ATM-Options

See Also: ATM-Direct-Profile

ATM-Direct-Profile

Description: Specifies the name of the Connection profile to which ATM data is Switched.

Usage: Specify a text string. The default is null.

Example: `set atm-direct-profile = myprof`

Dependencies: If ATM-Direct-Enabled is set to yes, you must specify a value for ATM-Direct-Profile.

Location: Connection > ATM-Options

See Also: ATM-Direct-Enabled

RADIUS attribute reference entries

Ascend-ATM-Direct (76)

Description: Specifies whether ATM-direct is enabled.

Usage: Specify ATM-Direct-Yes (1) to enable ATM-direct or ATM-Direct-No (0) to disable it.

See Also: Ascend-ATM-Direct-Profile (77)

Ascend-ATM-Direct-Profile (77)

Description: Specifies the hostname of the ATM interface to which data is switched.

Usage: Specify a text string. The default is null.

Dependencies: If Ascend-ATM-Direct is set to ATM-Direct-Yes, you must specify a value for Ascend-ATM-Direct-Profile.

See Also: Ascend-ATM-Direct (76)


ATM switched virtual circuits (SVCs)

With MAX TNT TAOS 8.0.0, the MAX TNT supports switched virtual circuit (SVC) services on DS-3 and OC-3 ATM interfaces. An *interface* is a point of ingress or egress to the system. An *ATM interface* is the logical configuration that enables ATM data to be sent and received.

An SVC is a point-to-point switched connection, which provides a lower-cost, usage-based alternative to ATM PVCs. Like other types of switched connections, SVCs can be initiated by a dial-in or a dial-out call, which can be made by the system on the basis of IP routing.


A dial-in ATM SVC terminates locally. The MAX TNT receives the call on an ATM interface. An example of a terminating SVC is shown in Figure 30.

Figure 30. Terminating SVC on an ATM interface



A dial-out ATM SVC is initiated as an outbound call on an ATM interface, either due to an explicit dial-out or on the basis of IP routing. Figure 31 shows a Pipeline unit dialing into the MAX TNT using PPP or some other type of encapsulation. The MAX TNT establishes the inbound call and then dials out on an ATM interface on the basis of IP routing, just as it would for another type of switched dial-out call.

Figure 31. Dial-out SVC on an ATM interface



Unlike permanent virtual circuits (PVCs), which require nailed connections, SVCs are on-demand connections and must use ATM endpoint addresses to identify the interface and route to it. To set up an SVC, you must configure SVC options, including an ATM address, in these locations:

- ATM-Interface profile, for a logical ATM interface associated with a physical ATM port
- Connection profile, used to establish the switched connection on an ATM interface

With the current software version, the system creates a static call route for an ATM address in each ATM-Interface profile. You can choose to configure the static call route explicitly using the ATMSVC-Route profile.

Current limitations

In this release, the ATM SVC implementation is subject to the following limitations:

- Because the Interim Local Management Interface (ILMI) is not implemented, dynamic address registration is not supported. Therefore, each ATM interface must be configured with a full SVC address.
- Only one ATM logical interface is supported for each ATM physical interface.

Address formats for ATM interfaces

The ATM endpoint address assigned to a MAX TNT ATM interface can be an ATM End System Address (AESAs) format or native E.164 address. AESA addresses are required for IP over ATM.

AESA formats

AESA addresses are 20-byte, 40-digit hexadecimal numbers. The first 13 bytes are the *address prefix*, or network portion of the address. The last 7 bytes are the host portion of the address.

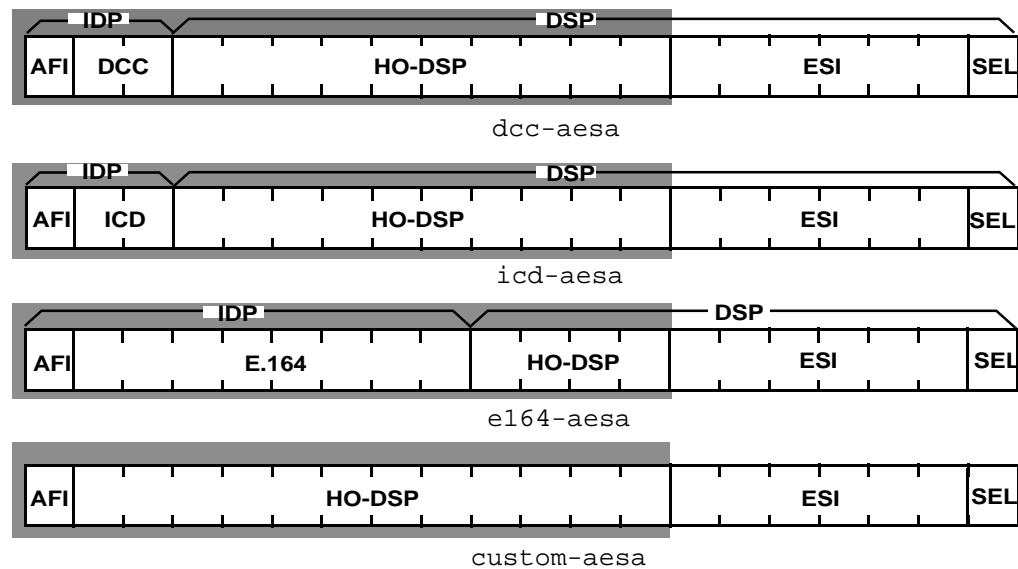
Depending on the AESA format chosen, the contents of each byte of the address varies, as shown in Figure 32. All of the supported AESA formats divide the address into the initial domain part (IDP), which defines the type of address and the regulatory authority responsible for allocating and assigning the domain-specific part, and the domain specific part (DSP).

AESA addresses use one of the following formats:

AESA format	Description
dcc-aesa	Data Country Code (DCC) is specified in the address, identifying the country in which the address is registered. Country codes are standardized and defined in ISO Reference 3166.
icd-aesa	International Code Designator (ICD) is specified in the address, identifying an international organization. The British Standards Organization administers these values.
e164-aesa	E.164 address is specified using the international format.
custom-aesa	Custom authority and format identifier (AFI) and byte order.

Figure 32 shows how each format divides the 20-byte address into subfields. The shaded portion represents the address prefix, which is always the first 13 bytes.

Figure 32. Subfields in the AESA address formats



For details about subfields in each format, see “Assigning an AESA format address” on page 217.

Native E.164 address format

Native E.164 addresses are regular ISDN numbers, including telephone numbers. E.164 addresses can contain up to 15 ASCII digits. For example, standard 10-digit United States telephone numbers, such as 508-555-1234, are native E.164 addresses.

Overview of configuring a physical ATM port

The DS3-ATM and OC3-ATM profiles require no special configuration to support SVCs. For information about configuring the physical ports, see the configuration guide that accompanied the slot card.

To access the configuration guides online, go to <http://www.ascend.com/doclibrary>. After you register, you can view or download the guides.

Overview of SVC options on a logical ATM interface

Following are the parameters (shown with default settings) for configuring a logical ATM SVC interface on a DS3-ATM or OC3-ATM port:

```
[in ATM-INTERFACE/{ { any-shelf any-slot 0 } 0 }]
interface-address* = { { any-shelf any-slot 0 } 0 }
name = ""

[in ATM-INTERFACE/{ { any-shelf any-slot 0 } 0 }:svc-options]
enabled = no
atm-protocol = uni-3.1
atm-address = { undefined "" { undefined { "" "" } { "" "" "" } } }
insert-calling-party-addr = yes
q93b-options = { 2 1 4000 30000 0 4000 4000 120000 4000 }
qsaal-options = { 64 4 25 67 1000 0 0 0 15000 }
```

Parameter	Specifies
Interface-Address	Interface address. This parameter includes the physical interface address (the shelf number, slot number, and item number of a port) and the logical-item number of the interface. Because only one ATM interface per physical ATM line is supported in this release, a logical-item value other than zero is not supported.
Name	Name of the ATM interface. The name can consist of up to 15 characters. The name is optional, and is used for informational purposes only.
Enabled	Enable/disable SVC signaling. If SVC signaling is enabled, a signaling PVC is created on the link to carry out SVC signaling and handle control messages. Q.93B and Q.SAAL layers are also initialized and enabled.
ATM-Protocol	ATM signaling protocol. The current implementation supports UNI 3.0 and UNI 3.1 protocols for SVCs. UNI 3.1 is selected as the factory default.
ATM-Address	AESA or E.164 address assigned to the interface.
Insert-Calling-Party-Addr	Enable/disable insertion of the calling-party address in outbound calls. If set to yes (the default), the system includes the calling party address on outbound calls. If set to no, the system does not include the calling party address on outbound calls.
Q93B-Options	Q.93B layer settings.
QSAAL-Options	Q.SAAL layer settings.

Assigning a native E.164 address

Following are the relevant parameters for assigning a native E.164 format ATM address, shown with default settings:

```
[in ATM-INTERFACE/{ { any-shelf any-slot 0 } 0 }:svc-options:atm-
address]
numbering-plan = undefined
```



```
e164-native-address = ""
svc-address-info = ""
```

Parameter	Specifies
Numbering-Plan	Type of SVC address. The default value is <code>undefined</code> , which indicates that an address has not been configured on the interface. To specify an E.164 address, set this parameter to <code>isdn</code> . To specify an AESA address, set it to <code>aesa</code> . The unknown and <code>x121</code> values are currently unsupported and have the same effect as the default <code>undefined</code> .
E164-Native-Address	SVC address in native E.164 format, up to 30 characters. For example, enter 5085552600 (a standard 10-digit U.S. telephone number).
AESA-Address	Does not apply to addresses in native E.164 format. See “Assigning an AESA format address” next.
SVC-Address-Info	Assigned address in read-only ASCII string format. For informational purposes only.

Assigning an AESA format address

The 20 bytes of an AESA address contain subfields, the size and contents of which may differ depending on the AESA format in use. The subfields are organized into an IDP portion and a DSP portion.

- The IDP portion specifies the authority and format identifier (AFI) and initial domain identifier (IDI) subfields.
- The DSP portion specifies the high-order domain-specific part (HO-DSP), end system identifier (ESI), and selector (SEL) subfields.

Note: The combination of IDP + HO-DSP + ESI must be unique. To ensure interoperability and equipment portability, use an ESI that is globally unique. For instance, if the ESI is not globally unique, and you move the ATM end system from one network to a different network, address conflicts can result.

For background information, see “AESA formats” on page 214. Following are the relevant parameters for assigning an AESA format address, shown with default settings:

```
[in ATM-INTERFACE/{ { any-shelf any-slot 0 } 0 }:svc-options:atm-
address]
numbering-plan = undefined
aesa-address = { undefined { "" "" } { "" "" "" } }
svc-address-info = ""

[in ATM-INTERFACE/{ { any-shelf any-slot 0 } 0 }:svc-options:atm-
address:aesa-address]
format = undefined

[in ATM-INTERFACE/{ { any-shelf any-slot 0 } 0 }:svc-options:atm-
address:aesa-address:idp-portion]
afi = ""
idi = ""

[in ATM-INTERFACE/{ { any-shelf any-slot 0 } 0 }:svc-options:atm-
address:aesa-address:dsp-portion]
ho-dsp = ""
```

```
esi = ""
sel = ""
```

Parameter	Specifies
Numbering-Plan	Type of SVC address. The default value is undefined, which indicates that an ATM address on the interface has not been configured. To specify an E.164 address, set this parameter to <code>isdn</code> . To specify an AESA address, set it to <code>aesa</code> . The unknown and <code>x121</code> values are currently unsupported and have the same effect as the default undefined.
SVC-Address-Info	Assigned address in read-only ASCII string format. For informational purposes only.
AESA-Address Format	AESA format for the interface. The default value is undefined, which indicates that the address has not been configured. Valid settings are <code>dcc-aesa</code> , <code>icd-aesa</code> , <code>e164-aesa</code> , and <code>custom-aesa</code> . For background information, see “AESA formats” on page 214.
IDP-Portion AFI	Hexadecimal code that identifies the kind of AESA address, such as DCC, ICD, or E.164 part of the AESA address, as well as the syntax of the rest of the address. The AFI is one byte, which contains two hex digits. For example, <code>0x39</code> (for <code>dcc-aesa</code>), <code>0x47</code> (for <code>icd-aesa</code>), or <code>0x45</code> (for <code>e164-aesa</code>).
IDP-Portion IDI	Hexadecimal code that identifies the subauthority that has allocated the address. For <code>dcc-aesa</code> and <code>icd-aesa</code> , the IDI is 2 bytes long (4 digits). For <code>e164-aesa</code> , the IDI is 8 bytes long, containing 16 digits that specify the E.164 address. The E.164 address can be up to 15 digits, so the system pads the number with leading zeros as required.
DSP-Portion HO-DSP	Hexadecimal number that specifies a segment of address space assigned to a particular device or network. For <code>dcc-aesa</code> and <code>icd-aesa</code> , the HO-DSP field is 10 bytes long, containing 20 hex digits. For <code>e164-aesa</code> , it is 4 bytes long (8 hex digits), and for <code>custom-aesa</code> it is 12 bytes long (24 hex digits).
DSP-Portion ESI	Hexadecimal number that uniquely identifies the end system within the specified subnetwork, typically an IEEE MAC address. This field is always 6 bytes long (12 hex digits).
DSP-Portion SEL	Hexadecimal number that is not used for ATM routing, but can be used by the end system. This subfield is always 1 byte long (2 hex digits).

Configuring the Q.93B layer

Q.93B parameters specify the timers as well as retry values associated with the functionality of the Q.93B layer. Following are the relevant parameters, shown with default settings:

```
[in ATM-INTERFACE/{ { shelf-1 slot-4 1 } 0 }:svc-options:q93b-options]
max-restart = 2
max-statenq = 1
t303-ms = 4000
t308-ms = 30000
t309-ms = 0
```

t310-ms = 4000
t313-ms = 4000
t316-ms = 120000
t322-ms = 4000

Parameter	Specifies
Max-Restart	Maximum number of unacknowledged transmit RESTART messages (from 1 to 32). The default value is 2.
Max-Statenq	Maximum number of unacknowledged transmit STATUS ENQ messages (from 1 to 32). The default value is 1.
T303-ms	Timer (in milliseconds) for a response after SETUP message is sent. The timer is stopped when a CONNECT, CALL PROCEEDING, or RELEASE COMPLETE message is received. Valid values are from 500 to 5000. The default value is 4000.
T308-ms	Timer (in milliseconds) for a response after a RELEASE message is sent. This is a <i>release indication timer</i> . The timer is started when the RELEASE message is sent and normally is stopped when the RELEASE or RELEASE COMPLETE message is received. Valid values are from 5000 to 50000. The default value is 30000.
T309-ms	Timer (in milliseconds) for Q.SAAL to reconnect. After this time has elapsed, calls are dropped. When set to 0 (the default), a default value based on an ATM signaling protocol is used. Valid values are from 0 to 200000.
T310-ms	Timer (in milliseconds) for a response after a SETUP message is received. This timer is also called the <i>call proceeding timer</i> . Valid values are from 5000 to 50000. The default value is 4000.
T313-ms	Timer (in milliseconds) for a response after a CONNECT message is sent. This timer is also called the <i>connect request timer</i> . The timer is started when the CONNECT message is sent and is stopped when the CONNECT ACKNOWLEDGE message is received. Valid values are from 1000 to 10000. The default value is 4000.
T316-ms	Timer (in milliseconds) for a response after a RESTART message is sent. This timer is also called the <i>restart request timer</i> . The timer is started when the RESTART message is sent and is stopped when the RESTART ACKNOWLEDGE message is received. Valid values are from 10000 to 300000. The default value is 120000.
T322-ms	Timer (in milliseconds) for a response after a STATUS ENQ message is sent. Valid values are from 1000 to 10000. The default value is 4000.

Configuring the Q.SAAL layer

Q.SAAL parameters specify the timers as well as retry values associated with the functionality of the Q.SAAL layer. Following are the relevant parameters, shown with default settings:

```
[in ATM-INTERFACE/{ { shelf-1 slot-4 1 } 0 }:svc-options:qsaal-
options]
window-size = 64
max-cc = 4
max-pd = 25
```

```

max-stat = 67
tcc-ms = 1000
tpoll-ms = 0
tkeepalive-ms = 0
tnoresponse-ms = 0
tidle-ms = 15000

```

Parameter	Specifies
Window-Size	Q.SAAL window size. Valid values are from 16 to 128. The default value is 64.
Max-Cc	Maximum number of control protocol data unit (PDU) retransmissions (BGN, END, RESYNC). Valid values are from 0 to 64. The default value is 4.
Max-PD	Maximum number of sequenced data PDFs between poll intervals. Valid values are from 1 to 64. The default value is 25.
Max-Stat	Maximum length of STAT PDU. Valid values are from 32 to 128. The default value is 67.
Tcc-ms	Retry time (in milliseconds) for control PDUs (BGN, END, RESYNC). Valid values are from 0 to 3000. The default value is 1000.
Tpoll-ms	Poll interval (in milliseconds) when active. When set to 0 (the default), a default value based on an ATM signaling protocol is used. Valid values are from 0 to 3000.
Tkeepalive-ms	Poll interval (in milliseconds) when in a transient state. When set to 0 (the default), a default value based on an ATM signaling protocol is used. Valid values are from 0 to 3000.
Tnoresponse-ms	Maximum interval (in milliseconds) between receipt of STAT PDUs. When set to 0 (the default), a default value based an ATM signaling protocol is used. Valid values are from 0 to 20000.
Tidle-ms	Poll interval (in milliseconds) when idle, for UNI 3.1 only. The default value is 15000. Valid values are from 1000 to 20000

Overview of SVC options in a Connection profile

The ATM options in a Connection profile are not specifically related to SVC configuration. VPI-VCI pairs are assigned by the switch for ATM SVCs. Most of the other settings in the ATM-Options subprofile operate in a similar manner for SVCs as they do for PVCs, once the SVC connection has been established.

Following are the parameters (shown with default settings) that are specific to configuring an ATM SVC connection:

```

[in CONNECTION/""]
encapsulation-protocol = atm
dial-number = ""

[in CONNECTION/"":ip-options]
ip-routing-enabled = yes
remote-address = 0.0.0.0/0

[in CONNECTION/"":atm-options:svc-options]
enabled = no

```

```
[in CONNECTION/"":atm-options:svc-options:incoming-caller-addr]
numbering-plan = undefined
e164-native-address = ""
aesa-address = { undefined { "" "" } { "" "" "" } }
svc-address-info = ""

[in CONNECTION/"":atm-options:svc-options:outgoing-called-addr]
numbering-plan = undefined
e164-native-address = ""
aesa-address = { undefined { "" "" } { "" "" "" } }
svc-address-info = ""
```

Parameter	Specifies
Encapsulation-Protocol	Encapsulation method for the connection. Must be set to atm for ATM SVC connections. Encapsulation-Protocol set to atm suggests that IP-over-ATM is used on the VC.
Dial-Number	Dial number for outbound calls. For dial-out ATM SVCs, you do not need to set this value. The system sets it to the same value as the outgoing-called-addr parameter when you write the Connection profile.
IP-Options IP-Routing-Enabled	Enable/disable IP routing for the interface. IP routing must be enabled (as it is by default) for outbound SVCs that are dialed on the basis of IP routing.
IP-Options Remote-Address	IP address of the far-end device, which can include a subnet specification. If it does not include a subnet mask, the router software in the MAX TNT unit assumes a default subnet mask that is based on address class.
SVC-Options Enabled	Enable/disable SVC for the connection. SVC is disabled by default.
SVC-Options Incoming-Caller-Addr	ATM address of the far end of the dial-in SVC connection, used to authenticate the inbound call. The address subfields operate in exactly the same way as the subfields of the same name in the ATM-Interface profile. For details, see “Assigning a native E.164 address” on page 216 or “Assigning an AESA format address” on page 217.
SVC-Options Outgoing-Called-Addr	ATM address of the far end of the dial-out SVC connection used to dial outbound SVC calls. The address subfields operate in exactly the same way as the subfields of the same name in the ATM-Interface profile. For details, see “Assigning a native E.164 address” on page 216 or “Assigning an AESA format address” on page 217.

Note: An SVC that can be initiated by either a dial-in or dial-out call specifies the same ATM address in both the incoming caller-addr and outgoing-called-addr fields.

Configuring a static ATM SVC route

With the current software version, no more than one ATM-Interface profile can be created for each physical ATM port, and the system creates an internal call route to the logical interface. As a result, you need not create explicit ATM static routes in this release. However, some sites

specify the route explicitly, to simplify route management. Following are the relevant parameters, shown with default settings, for creating an ATM static route:

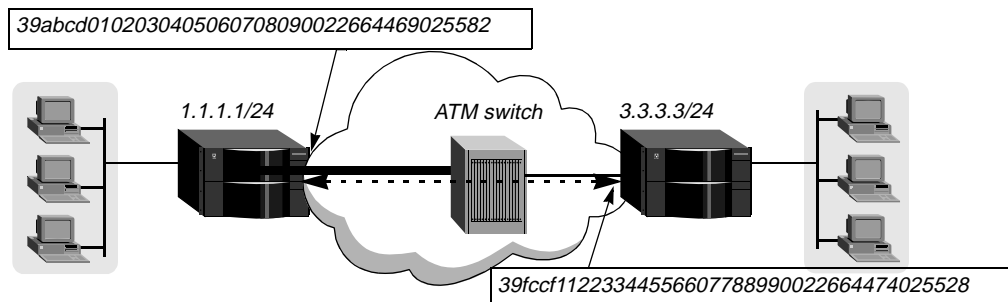
```
[in ATMSVC-ROUTE/" " ]
name* = ""
active = no
address-prefix = ""
interface-address = { { any-shelf any-slot 0 } 0 }
```

Parameter	Specifies
Name	Name of the route profile, up to 31 characters.
Active	Enable/disable the route for use.
Address-Prefix	The address prefix of the ATM address assigned to the interface in an ATM-Interface profile. For AESA-format addresses, the address prefix is the first 26 digits of the 40-digit hexadecimal number. For E.164 addresses, it is the entire address.
Interface-Address	Interface address of the ATM-Interface profile.

Example of configuring an ATM SVC

In the example shown in Figure 33, the system at the left side of the ATM cloud (the unit with the IP address 1.1.1.1/24) is a MAX TNT unit that will be configured for an ATM SVC that can be brought up by dial-in or dial-out on an OC3-ATM port. For the purposes of this example, the remote system (at the right side of the ATM cloud, with the IP address 3.3.3.3/24) is a DSLTNT unit, which has also been configured for an ATM SVC. The SVC can be brought up by a call to or from the remote system.

Figure 33. Example ATM SVC with DCC-AESA addresses



This example shows how to configure the MAX TNT unit at the left side of the ATM cloud (the unit with the IP address 1.1.1.1/24) in Figure 33.

Configuring the physical interface

The following commands configure the OC3-ATM physical interface and enable it for use:

```
admin> read oc3-atm { 1 2 1 }
OC3-ATM/{ shelf-1 slot-2 1 } read
admin> set name = atmswitch
admin> set enabled = yes
admin> set line-config clock-source = eligible
```

```
admin> write
OC3-ATM/{ shelf-1 slot-2 1 } written
```

Configuring the SVC logical interface

The following commands configure the SVC interface for the OC3-ATM port:

```
admin> read atm-interface { { 1 2 1 } 0 }
ATM-INTERFACE/{ { shelf-1 slot-2 1 } 0 } read

admin> set name = atmswitch

admin> set svc-options enabled = yes

admin> set svc-options atm-address numbering-plan = aesa

admin> list svc-options atm-address aesa-address
[in ATM-INTERFACE/{ { shelf-1 slot-2 1 } 0 }:svc-options:atm-address:+
format = undefined
idp-portion = { "" "" }
dsp-portion = { "" "" "" }

admin> set format = dcc-aesa

admin> set idp-portion afi = 39

admin> set idp-portion idi = abcd

admin> set dsp-portion ho-dsp = 01020304050607080900

admin> set dsp-portion esi = 226644690255

admin> set dsp-portion sel = 82

admin> write
ATM-INTERFACE/{ { shelf-1 slot-2 1 } 0 } written
```

Configuring a Connection profile to the far-end device

The following commands create a Connection profile to the far-end DSLTNT:

```
admin> new connection hanif-dsltn
CONNECTION/hanif-dsltn read

admin> set active = yes

admin> set encapsulation-protocol = atm

admin> set ip-options remote-address = 3.3.3.3/24

admin> set atm-options svc-options enabled = yes

admin> set atm-options svc incoming-caller-addr numbering-plan = aesa

admin> set atm-options svc outgoing-called-addr numbering-plan = aesa
```

In the following set of commands, notice that the incoming-caller-addr and outgoing-called-addr addresses are the same. This configuration allows the SVC to be brought up by a call to or from the far-end DSLTNT.

```
admin> list atm-options svc-options incoming-caller-addr aesa-address
[in CONNECTION/hanif-dsltn:atm-options:svc-options:incoming-caller-
addr:aesa-addres+
format = undefined
idp-portion = { "" "" }
dsp-portion = { "" "" "" }

admin> set format = dcc-aesa
```

```

admin> set idp-portion afi = 39
admin> set idp-portion idi = fccf
admin> set dsp-portion ho-dsp = 112233445566077889900
admin> set dsp-portion esi = 226644740255
admin> set dsp-portion sel = 28

admin> list .. .. outgoing-called-addr aesa-address
[in CONNECTION/hanif-dsltn:atm-options:svc-options:outgoing-called-
addr:aesa-address+
format = undefined
idp-portion = { " " " " }
dsp-portion = { " " " " " " }

admin> set format = dcc-aesa
admin> set idp-portion afi = 39
admin> set idp-portion idi = fccf
admin> set dsp-portion ho-dsp = 112233445566077889900
admin> set dsp-portion esi = 226644740255
admin> set dsp-portion sel = 28
admin> write
CONNECTION/hanif-dsltn written

```

When you write the profile with `outgoing-called-addr` configured, the system uses the configured value to set the `dial-number` parameter.

Command reference entry for new *ATMSVCroute* command

ATMSVCroute

Description: Displays the SVC call routing table. The system creates an SVC call routing entry for each configured ATM-Interface profile. To make an outbound call to a given destination ATM SVC address, the system consults the SVC call routing table for an address prefix. When the system finds a matching address prefix in the routing table, it uses the specified ATM-Interface profile index to route the call.

Permission level: System

Usage: `atmsvcroute [-d] | [-t]`

Option	Description
<code>-d</code>	Displays the SVC routing table.
<code>-t</code>	Toggles debug output.

Example: `atmsvcroute -d`

```
Prefix=39adfc01020304050507080900, lnk={{1, 4, 1}0}
```

The sample output show a single entry in the routing table. The first item in the entry is the address prefix of the destination ATM SVC address. The second item is the index of the ATM-Interface profile used to route the call.

ATM traffic shaping

With MAX TNT TAOS 8.0.0, you can control the rate of data transmission on an ATM interface. Each ATM interface supports up to 15 traffic shapers that define characteristics for different types of traffic. For example, voice traffic requires a constant amount of bandwidth and cannot tolerate delays, whereas file transfer can tolerate delay and variable bandwidth. Once you have specified the traffic shapers you need, you can apply a shaper to any number of connections.

Overview of ATM traffic-shaping settings

Following are the relevant parameters, shown with default values:

```
[in DS3-ATM/{ any-shelf any-slot 0 }:line-config:traffic-shapers[1]]
enabled = no
bit-rate = 1000
peak-rate = 1000
max-burst-size = 2
aggregate = no
priority = 0

[in OC3-ATM/{ any-shelf any-slot 0 }:line-config:traffic-shapers[1]]
enabled = no
bit-rate = 1000
peak-rate = 1000
max-burst-size = 2
aggregate = no
priority = 0

[in CONNECTION/":session-options]
traffic-shaper = 16
```


Parameter	Specifies
Enabled	Enable/disable the shaper for use.
Bit-Rate	Average bit rate in kilobits per second for transmitting traffic to the network. The default is 1000 (1 Mbps). For DS3-ATM interfaces, the valid range is from 0 to 37920. For OC3-ATM, the valid range is from 0 to 135631.
Peak-Rate	Maximum bit rate in kilobits per second for transmitting traffic to the network. The default is 1000 (1 Mbps). For DS3-ATM interfaces, the valid range is from 0 to 37920. For OC3-ATM, the valid range is from 0 to 135631.
Max-Burst-Size	Maximum burst size (MBS), which is the maximum number of cells that can be transmitted at Peak-Rate before the MAX TNT unit determines that the connection is exceeding the defined characteristics. The default is 2. The valid range is from 2 to 255.
Aggregate	Enable/disable aggregation of the Bit-Rate values of multiple VCs using this shaper. If set to no (the default), aggregation is not used. If set to yes, and <i>N</i> VCs are using this shaper, the throughput of each VC is Bit-Rate/ <i>N</i> .
Priority	Priority of this shaper relative to other shapers on this interface. The valid range is from zero (the default) to 15. Zero indicates the highest priority, and 15 indicates the lowest.

Parameter	Specifies
Traffic-Shaper	The traffic shaper assigned to the connection. The default is shaper 16, which is an internal shaper that is not configurable.

Example of configuring traffic shaping

In the example shown in Figure 34, the MAX TNT unit has a DS3-ATM interface to a DSL Terminator unit, a Frame Relay data link interface to a Pipeline 220 unit, and an ATM-Frame Relay circuit between the two interfaces.

Figure 34. Example traffic shaping setup



The following commands define a data link to the Pipeline 220 on a nailed T1 line (nailed group 999), which a bit rate of approximately 1.5 Mbps :

```
admin> new frame ut1-p220
FRAME-RELAY/ut1-p220 read
admin> set active = yes
admin> set nailed-up-group = 999
admin> set link-type = nni
admin> write
FRAME-RELAY/ut1-p220 written
```

The following commands configure the DS3-ATM interface and define a traffic shaper that limits the bit rate to less than 500 Kbps:

```
admin> read ds3-atm {1 3 1}
DS3-ATM/{ shelf-1 slot-3 1 } read
admin> set name = atm-switch
admin> set enabled = yes
admin> set line nailed-group = 111
admin> set line high-tx-output = yes
admin> set line traffic-shapers 1 enabled = yes
admin> set line traffic-shapers 1 bit-rate = 500
admin> write
ATM-DS3/{ shelf-1 slot-3 1 } written
```

The following commands specify the circuit between the Frame-Relay and ATM interfaces, and apply the traffic shaper to the ATM interface:

```
admin> new conn p220
CONNECTION/p220 read
```

```
admin> set active = yes
admin> set encapsulation-protocol = frame-relay-circuit
admin> set ip-options ip-routing-enabled = no
admin> set telco-options call-type = ft1
admin> set fr-options frame-relay-profile = fr-switch
admin> set fr-options dlci = 100
admin> set fr-options circuit-name = atmfr-1
admin> write
CONNECTION/p220 written
admin> new conn terminator
CONNECTION/terminator read
admin> set active = yes
admin> set encapsulation-protocol = atm-frame-relay-circuit
admin> set ip-options ip-routing-enabled = no
admin> set fr-options circuit-name = atmfr-1
admin> set telco-options call-type = ft1
admin> set telco-options nailed-groups = 111
admin> set session-options traffic-shaper = 1
admin> set atm-options vpi = 100
admin> set atm-options vci = 132
admin> write
CONNECTION/terminator written
```

The traffic shaper in the DS3-ATM profile does not enable aggregation (the default setting), so the actual transfer rate across the connection to the Terminator will be approximately 480 Kbps, as the shaper permits.

If two VCs were configured on the DS3-ATM interface, both using a shaper that specified a bit rate of 500 with aggregate set to yes, each VC would use a transfer rate of about half the specified bit rate, or 240 Kbps.

ATM-Frame Relay transparent-mode circuits (FRF.8)

In earlier releases, the MAX TNT supported translation-mode ATM-Frame Relay circuits. In translation mode, the system removes Frame Relay Multiprotocol Encapsulation (RFC 1490) from the data stream received on a Frame Relay interface, and adds ATM Multiprotocol Encapsulation (RFC 1483) to the data stream sent on an ATM interface, and vice versa, from one side of the circuit to the other.

With MAX TNT TAOS 8.0.0, the MAX TNT supports transparent-mode ATM-Frame Relay circuits, as defined in the *FRF.8 Frame Relay ATM/PVC Service Interworking Implementation Agreement*. In transparent mode, the system performs no conversion, but simply passes the data stream from one side of the circuit to the other.

Transparent mode is supported on DS3-ATM and OC3-ATM cards in this release. Transparent mode requires that the circuit endpoints support compatible upper-layer protocols for applications such as packetized voice.

Overview of Connection profile settings

Following is the relevant parameter, shown with its default value:

```
[in CONNECTION/"":atm-options]
fr-08-mode = translation
```

Parameter	Specifies
FR-08-Mode	Translation or transparent mode of operation for the ATM-Frame Relay circuit. The default is translation mode, which causes the system to convert RFC 1490 encapsulation to RFC 1483, and vice versa. In transparent mode, the data is passes from one side of the circuit to the other without 1490-to-1483 translation. The encapsulation mode for the profile must be atm-frame-relay-circuit for this parameter to have an effect.

Overview of RADIUS profile settings


RADIUS uses the following attribute-value pair to specify a transparent mode circuit:

RADIUS attribute	Value
Ascend-FR-08-Mode (30)	Translation (0) or transparent (1) mode of operation for the ATM-Frame Relay circuit. The default is translation mode, which causes the system to convert RFC 1490 encapsulation to RFC 1483, and vice versa. In transparent mode, the data is passes from one side of the circuit to the other without 1490-to-1483 translation. The encapsulation mode for the profile must be atm-frame-relay-circuit for this setting to have an effect.

Example of configuring a transparent-mode ATM-Frame Relay circuit

In the example shown in Figure 35, the MAX TNT receives frames on a Frame Relay DLCI interface and transmits them on an ATM PVC (and vice versa) without removing the frames' encapsulation and adding the encapsulation required by the other endpoint.

Figure 35. ATM-Frame Relay circuit



Using local profiles

The following commands define the data link to the Frame Relay switch:

```
admin> new frame fr-switch
FRAME-RELAY/fr-switch read

admin> set active = yes

admin> set nailed-up-group = 999
```

```
admin> write
FRAME-RELAY/fr-switch written
```

The following commands configure a DS3-ATM interface:

```
admin> read ds3-atm {1 3 1}
DS3-ATM/{ shelf-1 slot-3 1 } read
admin> set name = atm-switch
admin> set enabled = yes
admin> set line nailed-group = 111
admin> set line high-tx-output = yes
admin> write
ATM-DS3/{ shelf-1 slot-3 1 } written
```

The following commands specify the transparent mode circuit between the Frame Relay and ATM interfaces:

```
admin> new conn fr-endpoint
CONNECTION/fr-endpoint read
admin> set active = yes
admin> set encapsulation-protocol = frame-relay-circuit
admin> set ip-options ip-routing-enabled = no
admin> set telco-options call-type = ft1
admin> set fr-options frame-relay-profile = fr-switch
admin> set fr-options dlci = 100
admin> set fr-options circuit-name = atmfr-1
admin> write
CONNECTION/fr-endpoint written
admin> new conn atm-endpoint
CONNECTION/atm-endpoint read
admin> set active = yes
admin> set encapsulation-protocol = atm-frame-relay-circuit
admin> set ip-options ip-routing-enabled = no
admin> set fr-options circuit-name = atmfr-1
admin> set telco-options call-type = ft1
admin> set telco-options nailed-groups = 111
admin> set atm-options vpi = 100
admin> set atm-options vci = 132
admin> set atm-options fr-08-mode = transparent
admin> write
CONNECTION/atm-endpoint written
```

Using RADIUS profiles

The following frdlink pseudo-user profile defines the data link to the Frame Relay switch:

```
frdlink-sys-1 Password = "ascend"
  Service-Type = Dialout-Framed-User,
```

```
Ascend-FR-Profile-Name = "fr-switch",
Ascend-Call-Type = Nailed,
Ascend-FR-Type = Ascend-FR-NNI,
Ascend-FR-Nailed-Grp = 999
```

The DS3-ATM or OC3-ATM interface is configured in a local profile, as shown in the preceding section. The next set of profiles specifies the circuit between the Frame Relay and ATM interfaces:


```
permconn-sys-1 Password = "ascend"
  Service-Type = Dialout-Framed-User,
  User-Name = "fr-endpoint",
  Framed-Protocol = FR-CIR,
  Ascend-Route-IP = Route-IP-No,
  Ascend-FR-DLCI = 100,
  Ascend-FR-Profile-Name = "fr-switch",
  Ascend-FR-Circuit-Name = "atmfr-1"

permconn-sys-2 Password = "ascend"
  Service-Type = Dialout-Framed-User,
  User-Name = "atm-endpoint",
  Framed-Protocol = ATM-FR-CIR,
  Ascend-Route-IP = Route-IP-No,
  Ascend-Group = "111",
  Ascend-ATM-Vpi = 100,
  Ascend-ATM-Vci = 132,
  Ascend-FR-Circuit-Name = "atmfr-1",
  Ascend-FR-08-Mode = 1
```

ATM-Frame Relay virtual channel trunking

In previous releases, Frame Relay and ATM-Frame Relay circuits always had two endpoints (1:1 circuits). With MAX TNT TAOS 8.0.0, the MAX TNT also supports virtual channel trunking, which allows $N:1$ circuits. With virtual channel trunking, a circuit can have more than two endpoints, as long as multiple endpoints are designated as host links and only one endpoint is designated as a trunk link. The system aggregates traffic from multiple host links onto one trunk link, creating an $N:1$ circuit, as shown in Figure 36:

Figure 36. $N:1$ circuit between multiple Frame Relay hosts and an ATM trunk



With virtual channel trunking, the circuit endpoints can include multiple Frame Relay DLCI interfaces and an ATM VPI-VCI interface, as long as only one trunk link is specified.

When the system receives upstream traffic from a host link, it learns the host's MAC address and then forwards the data to the trunk-link interface. When the system receives downstream

traffic from the trunk link, it uses the destination MAC address to transmit the packets on the appropriate host link.

Current limitations

In this release, the virtual channel trunking implementation is subject to the following limitations:

- Only one ATM endpoint can be defined per circuit.
- Broadcast and multicast packets from the trunk link are not forwarded to the host links of the circuit.
- Packets from the individual host links are not forwarded to the other host links.

Overview of Connection profile settings

Following is the relevant parameter, shown with its default value:


```
[in CONNECTION/"":fr-options]
circuit-name = ""
fr-link-type = transparent-link
```

Parameter	Specifies
Circuit-Name	Circuit name (up to 16 characters). The other endpoint(s) of a circuit must specify the same circuit name.
FR-Link-Type	Type of link for the circuit endpoint. Valid values are <code>transparent-link</code> (the default), <code>host-link</code> , and <code>trunk-link</code> . A transparent-link circuit is a 1:1 circuit. It requires two endpoints that specify the same circuit name and the <code>transparent-link</code> type. If only one endpoint is specified, data received on the specified DLCI is dropped. If more than two transparent-link endpoints are specified with the same circuit name, only two of the profiles are used to form a circuit. Virtual channel trunking allows an N:1 circuit. It can have more than two endpoints that specify the same circuit name, as long as multiple endpoints specify the <code>host-link</code> type and only one endpoint specifies the <code>trunk-link</code> type.

Example of configuring virtual channel trunking

In the following example, two Frame Relay hosts are switched to an ATM trunk link, as shown in Figure 37.

Figure 37. Circuit using virtual channel trunking



The example commands do not include data link or physical link configurations. For details on those topics, see the *MAX TNT Network Configuration Guide* and *MAX TNT Hardware Installation Guide*.

The following commands configure the Connection profile for the ATM trunk link, where the nailed group is configured on an ATM interface:

```
admin> new conn atm-trunk1
CONNECTION/atm-trunk1 read

admin> set active = yes

admin> set encapsulation-protocol = atm-frame-relay-circuit

admin> set telco-options call-type = ft1

admin> set telco-options nailed-groups = 111

admin> set ip-options ip-routing-enabled = no

admin> set fr-options circuit-name = vtrunk-cir1

admin> set fr-options fr-link-type = trunk-link

admin> set atm-options vpi = 101

admin> set atm-options vci = 156

admin> write
CONNECTION/atm-trunk1 written
```

The following commands configure the Connection profile for the first Frame Relay host link:

```
admin> read conn frhost-1
CONNECTION/frhost-1 read

admin> set active = yes

admin> set encapsulation-protocol = frame-relay-circuit

admin> set ip-options ip-routing-enabled = no

admin> set telco-options call-type = ft1

admin> set fr-options frame-relay-profile = ct1.8-fr

admin> set fr-options dlci = 16

admin> set fr-options circuit-name = vtrunk-cir1

admin> set fr-options fr-link-type = host-link

admin> write
CONNECTION/frhost-1 written
```


The following commands configure the Connection profile for the second Frame Relay host link:

```
admin> read conn frhost-2
CONNECTION/frhost-2 read

admin> set active = yes

admin> set encapsulation-protocol = frame-relay-circuit

admin> set ip-options ip-routing-enabled = no

admin> set telco-options call-type = ft1

admin> set fr-options frame-relay-profile = ut1.3-fr

admin> set fr-options dlci = 100

admin> set fr-options circuit-name = vtrunk-cir1

admin> set fr-options fr-link-type = host-link

admin> write
CONNECTION/frhost-2 written
```

Signaling System 7 (SS7)

Support for Signaling System 7 (SS7) was introduced with limited availability in earlier TAOS 7.x releases. It is now generally available in MAX TNT TAOS 8.0.0. SS7 is an internationally standardized general-purpose common-channel signaling system designed for use over a variety of digital circuit-switched networks. At the physical layer, it uses T1, T3, or E1 for data traffic and separate TDM circuits for signaling information.

In MAX TNT TAOS 8.0.0, the following two methods of integration with an SS7 network are supported, each of which requires a separate software license:

- Access SS7 Gateway Control Protocol (ASGCP). This method of integration enables the MAX TNT to terminate data calls in an SS7 network. The signaling gateway must be ICD for softswitch (formerly ASG). ICD stands for Internet Call Diversion.
- IP Device Control (IPDC). IPDC is a third-party proprietary protocol. This method of integration enables the MAX TNT to terminate both voice and data calls. The signaling gateway can be ICD for softswitch or Lucent Softswitch.

Table 16 shows the protocols supported by these signaling gateway platforms.

Table 16. Signaling gateway platforms and protocol support

Platform	IPDC 0.12	ASGCP (Q.931+)
ICD for softswitch (formerly ASG)	Supported	Supported
Lucent Softswitch	Supported	Not supported

System requirements for SS7 operations

A MAX TNT unit configured for SS7 in communication with an SS7 signaling gateway is a service switching point (SSP). To operate in this capacity, the MAX TNT must have the following equipment and licenses:

- SS7 software license, either for ASGCP or IPDC


- Sufficient T1, T3, or E1 trunks
- Sufficient modem or Hybrid Access cards (or both) to terminate data calls
- One or more Ethernet cards (recommended to offload the shelf controller)

If the system will operate as a MultiVoice™ gateway in an SS7 environment, a MultiVoice software license must also be enabled and one or more MultiDSP cards must be installed to enable the system to terminate voice calls. For details about MultiVoice, see “MultiVoice operations” on page 255.

MAX TNT as terminator of data calls in an SS7 network

With the ASGCP license, MAX TNT units can decrease congestion on the Public Switched Telephone Network (PSTN) caused by users connecting to the Internet. An example of MAX TNT units being used for this purpose is shown in Figure 38.

Figure 38. MAX TNT terminating data calls in an SS7 network

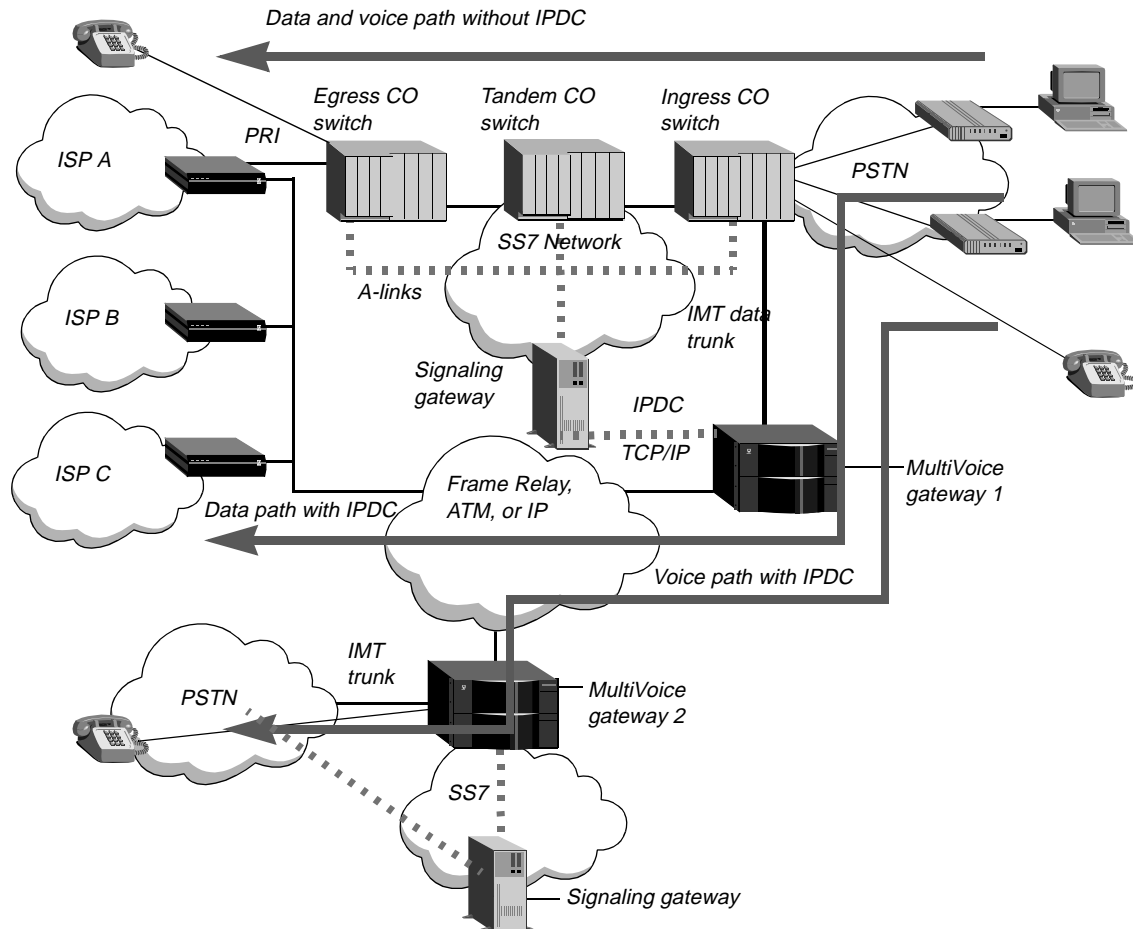


The MAX TNT is connected to the entry (ingress) central office (CO) switch via intermachine trunks (IMTs) and to a signaling gateway by means of dual-link (primary and secondary) TCP/IP links. Each CO switch is a service switching point (SSP). The combination of a MAX TNT and signaling gateway is also an SSP. The signaling gateway is connected to the SS7 network by access links (A-links). The signaling gateway and the MAX TNT together act as a switch that routes calls intended for ISPs directly to the MAX TNT, thus avoiding the PSTN tandem or transit switches and interoffice trunks.

MAX TNT as terminator of voice and data calls in an SS7 network

With the IPDC license, IPDC is used for communication between the signaling gateway and the MAX TNT. IPDC enables the MAX TNT to terminate voice or data calls. An example of MAX TNT units being used both for Internet call diversion (data) and Voice over IP (VoIP) is shown in Figure 39.

Figure 39. MAX TNT terminating voice and data calls in an SS7 network



Connection to the SS7 network is achieved through a signaling gateway. This gateway provides a bridge to the SS7 network and performs service switching point functions such as initiating and managing call setup and release, and executing call routing. IPDC must be supported by both the signaling gateway and the MAX TNT.

The signaling gateway uses the IPDC protocol to convert the SS7 signaling information and call data from the PSTN into IPDC packets, which are sent to the MAX TNT. In addition, the gateway uses IPDC to convert IPDC packets received from a MAX TNT into SS7 format before sending the call to the PSTN.

Before sending call data across the IP network, the MAX TNT uses IPDC to extract TDM and IP routing instructions from the IPDC packets received from the signaling gateway. The far-end MAX TNT then forwards IPDC packets to a signaling gateway, which converts them back into SS7 messages before the call is connected.

See “MultiVoice operations” on page 255 for a more detailed description of how VoIP calls are processed by IPDC.

Interface between a signaling gateway and MAX TNT

TCP/IP is the transport service used to carry control messages between a signaling gateway and the MAX TNT. The data delivery layer (DDL) uses a TCP/IP socket on both the signaling gateway and MAX TNT. On the signaling gateway side, the DDL is the server that listens for the socket connection and keeps track of the mapping between a MAX TNT unit and its socket. On the MAX TNT side, the DDL is the client that initiates a socket connection and handles connection establishment, connection recovery, and link selection.

Incoming calls

The ingress central office (CO) switch (see Figure 38) processes the incoming call based on the called number, then identifies the MAX TNT as the destination for the call. The SS7 network sends an initial address message (IAM) to the signaling gateway. The signaling gateway informs the MAX TNT that a call will be coming in on one of the IMT channels from the CO switch. The message from the CO switch contains the calling and called party number, the circuit identification code (CIC), and the destination point code (DPC). The signaling gateway sends an address complete message (ACM) to the SS7 network acknowledging that it has received the relevant information to route the call.

The signaling gateway then sends a call origination message to the MAX TNT to establish a path between the ingress switch and the MAX TNT. The MAX TNT sets up the path and then sends an answer message to the signaling gateway so that the signaling gateway can make the proper updates to its resource management database. For a T1 or T3 network, the signaling gateway then sends an answer message to the SS7 network.

Once the path is set up, the MAX TNT accepts the call, off-loading the Internet call from the PSTN to the data network. The data network used to off-load the call can be a Frame Relay, ATM, or IP network.

Continuity tests

A continuity test can be performed at the time of call setup or during testing to verify that the physical link between the CO switch and the MAX TNT is available. The CO switch informs the signaling gateway, which then informs the MAX TNT that it will conduct a continuity test on the circuit. During a call continuity test, the CO switch sends a tone through the physical path to the MAX TNT and receives a tone back from the MAX TNT indicating the continuity of the path.

SS7-Gateway profile settings

The signaling gateway and MAX TNT communicate over a TCP/IP link. The signaling interface can be a single or dual TCP connection between the MAX TNT and signaling gateway. When the interface initializes, it opens TCP connections to the specified addresses and ports of the signaling gateway. The MAX TNT keeps the TCP connections open as long as the unit is operating and the signaling interface is enabled.

Settings in the SS7-Gateway profile configure the signaling interface. The MAX TNT resets the signaling link whenever changes are written to the profile.

Following are the parameters (shown with default settings) for configuring the signaling interface:

```
[in SS7-GATEWAY]
enabled = no
control-protocol = asgcp
primary-ip-address = 0.0.0.0
primary-tcp-port = 0
secondary-ip-address = 0.0.0.0
secondary-tcp-port = 0
bay-id = ""
system-type = IASCTNT1B
transport-options = { 0 1000 3000 30000 7 6 no }
use-system-ip-address-as-source = yes
```

Parameter	Specifies
Enabled	Enable/disable the interface. When set to <code>no</code> (the default), the interface is disabled. When set to <code>yes</code> , the interface is enabled if the Primary-IP-Address and Primary-TCP-Port also have valid values. Changing the setting from <code>yes</code> to <code>no</code> closes the signaling links but does not disconnect active SS7 calls.
Control-Protocol	Control protocol. The <code>asgcp</code> setting enables the unit to terminate data calls by using ASGCP. The <code>ipdc-0.x</code> (XCOM/Level 3 IPDC) setting enables the unit to terminate voice and data using IPDC. If only one SS7 license is enabled, the parameter defaults to that control protocol (<code>asgcp</code> or <code>ipdc-0.x</code>) and cannot be modified. If both licenses are enabled, the parameter defaults to <code>asgcp</code> . See “Specifying the SS7 control protocol” on page 238 for more information about this parameter.
Primary-IP-Address Primary-TCP-Port	IP address and TCP port to use for communication with the primary signaling gateway. These settings are required for SS7 operations.
Secondary-IP-Address Secondary-TCP-Port	IP address and TCP port to use for communication with a secondary signaling gateway. These settings are optional. If specified, the secondary signaling gateway is used only when the primary gateway is unavailable. The primary and secondary address and port configurations can point to two Ethernet interfaces of the same signaling gateway.
Bay-ID	This parameter does not apply when Control-Protocol is set to <code>asgcp</code> . When Control-Protocol is set to <code>ipdc-0.x</code> , the system sends its value as an ASCII string to the media gateway controller in the device registration message. The MAX TNT does not interpret the value. Interpretation on the signaling gateway is gateway dependent.
System-Type	This parameter does not apply when Control-Protocol is set to <code>asgcp</code> . When Control-Protocol is set to <code>ipdc-0.x</code> , the system sends its value as an ASCII string to the media gateway controller in the device registration message. The MAX TNT does not interpret the value. Interpretation on the signaling gateway is gateway dependent.

Parameter	Specifies
Transport-Options	The Transport-Options subprofile contains settings for changing the operation of SS7 DDL timers. See “Configuring transport-layer options” on page 238.
Use-System-IP-Address-As-Source	Enable/disable use of the system address as the source address for packets generated by the MAX TNT unit. See “System IP address considerations” on page 239.

Specifying the SS7 control protocol

With the appropriate software license, the MAX TNT supports either ASGCP and IPDC 0.12 control protocol. If only one of the possible control protocols (`asgcp` or `ipdc-0.X`) is licensed on the MAX TNT, the Control-Protocol parameter defaults to the licensed protocol and cannot be modified. However, if both protocols are licensed, the parameter defaults to `asgcp`. Because of this default and because the MAX TNT does not store unmodified profile items in NVRAM, the setting can be modified unintentionally when you upgrade to new software or enable a new license to support a second control protocol. For this reason, Lucent recommends that you verify the setting after upgrading. If the proper protocol is not specified, change the setting and then reset the unit.

Although the control protocol is configurable in real time, you must reset the system to begin using the new protocol. After the MAX TNT is reset, it establishes a new TCP link to the signaling gateway and begins communicating with it using the specified control protocol.

Configuring transport-layer options

Administrators occasionally need to change the duration of various SS7 DDL timers to fine-tune a signaling link. For example, you might want to change timeouts when integrating a MAX TNT unit with existing signaling gateways. The following parameters, shown with default values, are used to set MAX TNT time intervals for waiting and responding to the various signaling link processes:

```
[in SS7-GATEWAY:transport-options]
device-id = 0
t1-duration = 1000
t2-duration = 3000
t3-duration = 30000
window-size = 7
ack-threshold = 6
heart-beat = no
```

Parameter	Specifies
Device-ID	Logical SS7 command control device where these values apply. Currently, the settings in this profile apply only to the MAX TNT unit's operations. <i>This parameter is currently not used.</i>
T1-Duration	Value of the acknowledgement (ACK) delay timer in milliseconds. This timer specifies the maximum delay for an acknowledgement when an information frame (I-frame) is received. The default value is 1000 (1 second). The value must be less than the T2 duration timer specified on the signaling gateway. Valid values range from 0 to 2147483647.

Parameter	Specifies
T2-Duration	Value of the transmission time-out timer in milliseconds. This timer specifies how long this endpoint must wait for an acknowledgement to a heartbeat frame. The default value is 3000 (3 seconds). The value must be greater than the T1 duration timer on the signaling gateway. Valid values range from 0 to 2147483647.
T3-Duration	Value of the persistent error timer in milliseconds. This timer specifies the maximum duration of attempts to reestablish a link before the transport layer flushes the data queues and sends an error indication up. Default value is 30000 (30 seconds). Valid values range from 0 to 2147483647.
Window-Size	Maximum number of sequentially numbered data packets that can be sent while pending acknowledgement at any given time. Default value is 7. Valid values range from 1 to 63.
Ack-Threshold	Threshold for triggering an acknowledgement (ACK) while receiving data packets. As soon as the specified number of packets is received, the MAX TNT sends an ACK back regardless of the value of its T1 timer. The value of this parameter must not be greater than the window size. Default value is 6. Valid values range from 1 to 63.
Heartbeat	Enable/disable detection of a physical link failure, such as disconnection of a cable or failure of the signaling gateway. When the parameter is set to <i>yes</i> , the MAX TNT periodically sends out heartbeat frames to the signaling gateway and waits for an acknowledgement. If it does not receive an acknowledgement within the number of milliseconds specified in its T2-Duration timer, the MAX TNT resets the signaling link.

System IP address considerations

The System-IP-Addr parameter of the IP-Global profile specifies the source address of all packets generated by the system, such as the connection request packets sent to a signaling gateway to establish communication. When the Use-System-IP-Address-As-Source parameter is set to *yes* (the default), the MAX TNT uses the system address as its source address in the packets it sends to the signaling gateway.

For some sites, administrative policy or other constraints introduce a requirement to use the system address for some purposes, but to use a separate source address for communication with the signaling gateway. For example, although a site might require a certain system address for compatibility with other routers, this requirement might cause an address space conflict, or might cause delays and time-outs in the receipt of acknowledgements from signaling gateways. Or, a site might decide to separate the signaling control network from the Internet for security purposes.

To enable sites to integrate MAX TNT units into their infrastructure and at the same time communicate efficiently with signaling gateways, the following parameter (shown with its default value) was introduced:

```
[in SS7-GATEWAY]
use-system-ip-address-as-source = yes
```

When this parameter is set to `no`, the MAX TNT does not use the system address as its source address for signaling packets. Instead, it uses the IP address of the Ethernet interface on which the signaling packets are sent. When the parameter is set to `yes`, the MAX TNT uses the same system address for signaling packets as for all other packets generated by the system.

Example of a basic configuration

The following commands configure an SS7-Gateway profile for a single TCP connection to a signaling gateway running IPDC:

```
admin> read ss7-gateway
SS7-GATEWAY read

admin> set enabled = yes

admin> set primary-ip-address = 1.1.1.1

admin> set primary-tcp-port = 5000

admin> write
SS7-GATEWAY written
```

Note: For the link to become active, the signaling gateway must have a matching entry for the MAX TNT. For information about configuring the signaling gateway, see the documentation that came with the unit.

T1 lines as SS7 data trunks

To configure T1 lines for SS7, you must set the following parameters, shown with sample settings:

```
[in T1/{ shelf-1 slot-1 7 }:line-interface]
signaling-mode = ss7-data-trunk
incoming-call-handling = internal-processing

[in T1/{ shelf-1 slot-1 7 }:line-interface:channel-config:24]
channel-usage = switched-channel
```

Parameter	Usage for SS7 data trunks
Signaling-Mode	<p>For an SS7 data trunk, which carries no signaling, this parameter can be set to either of the following values. The setting registers the line with the signaling gateway and allows the gateway to take control of the line and its calls.</p> <p><code>ss7-data-trunk</code> causes the unit to provide clear 64Kbps SS7 data trunk support. If any of the PSTN switches you are using is a 1AESS switch, which uses robbed-bit signaling, this setting can sometimes cause that switch to receive fluctuating A/B bit status. This condition might ultimately force the line out of service, unless you disable robbed-bit signaling on the 1AESS switch.</p> <p><code>ss7-robbed-bit</code> causes the MAX TNT to send a steady A/B bit status on the SS7 data trunk, which eliminates the need to disable robbed-bit signaling on the 1AESS switch.</p>

Parameter	Usage for SS7 data trunks
Incoming-Call-Handling	Specifies how the MAX TNT processes incoming calls on this line. For SS7 data trunks, the parameter must be set to <code>internal-processing</code> in this release. The <code>ss7-gateway-processing</code> setting for passing incoming call requests to an external signaling gateway is currently not supported.
Channel-Usage	T1 lines typically use channel 24 for signaling. For SS7 data trunks, the Channel-Usage setting for channel 24 should be <code>switched-channel</code> .

Example of configuring a T3 card for SS7 data

To configure lines of a T3 card as SS7 data trunks, you must first configure the T3 profile as in the following example:

```
admin> read t3 {1 1 1}
T3/{ shelf-1 slot-1 1 } read
admin> set enabled = yes
admin> set frame-type = m13
admin> set line-length = 0-225
admin> write
T3/{ shelf-1 slot-1 1 } written
```

After configuring the T3 line, configure the individual T1 lines that constitute the T3 line as explained in the next section.

Example of configuring a T1 data trunk

The following commands configure a T1 line as an SS7 data trunk, enabling the signaling gateway to control the line:

```
admin> read t1 {1 1 7}
T1/{ shelf-1 slot-1 7 } read
admin> set line-interface enabled = yes
admin> set line-interface signaling-mode = ss7-data-trunk
admin> set line-interface incoming-call-handling = internal-processing
admin> set line-interface channel-config 24 channel-usage = switched
admin> write
T1/{ shelf-1 slot-1 7 } written
```

E1 lines as SS7 data trunks

In MAX TNT TAOS 8.0.0, the MAX TNT supports E1 SS7 data trunks. Configuring the E1 SS7 data trunks is very similar to configuring T1 data trunks. To configure E1 lines for SS7, you must set the following parameters in an E1 profile, shown with sample settings:

```
[in E1/{ shelf-1 slot-10 1 }:line-interface]
signaling-mode = ss7-data-trunk
incoming-call-handling = internal-processing

[in E1/{ shelf-1 slot-10 1 }:line-interface:channel-config[17]]
channel-usage = switched-channel]
```

Parameter	Usage for SS7 data trunks
Signaling-Mode	<p>For an SS7 data trunk, which carries no signaling, this parameter can be set to either of the following values. The setting registers the line with the signaling gateway and allows the gateway to take control of the line and its calls.</p> <p><code>ss7-data-trunk</code> causes the MAX TNT to provide clear 64Kbps SS7 data trunk support. If any of the PSTN switches you are using is a 1AESS switch, which uses robbed-bit signaling, this setting can sometimes cause that switch to receive fluctuating A/B bit status. This condition might ultimately force the line out of service, unless you disable robbed-bit signaling on the 1AESS switch.</p> <p><code>ss7-robbed-bit</code> causes the MAX TNT to send a steady A/B bit status on the SS7 data trunk, which eliminates the need to disable robbed-bit signaling on the 1AESS switch.</p>
Incoming-Call-Handling	<p>Specifies how the MAX TNT processes incoming calls on this line. For SS7 data trunks, the parameter must be set to <code>internal-processing</code> in this release. The <code>ss7-gateway-processing</code> setting for passing incoming call requests to an external signaling gateway is currently not supported.</p>
Channel-Usage	<p>In the MAX TNT, the channel-config index begins with 1 (not 0), so E1 lines typically use channel 17 for signaling. For SS7 data trunks, change the default Channel-Usage setting for channel 17 from <code>d-channel</code> to <code>switched-channel</code>.</p>

For example, the following commands configure an E1 line as an SS7 data trunk, enabling the signaling gateway to control the line:

```
admin> read e1 {1 10 1}
E1/{ shelf-1 slot-10 1 } read

admin> set line-interface enabled = yes

admin> set line-interface signaling-mode = ss7-data-trunk

admin> set line-interface incoming-call-handling = internal-processing

admin> set line-interface channel-config 17 channel-usage = switched

admin> write
E1/{ shelf-1 slot-10 1 } written
```

V.110 bearer capability for SS7 calls using IPDC

With MAX TNT TAOS 8.0.0, the MAX TNT supports V.110 bearer capability for SS7 calls using IPDC. This feature enables SS7 call routing across V.110 interfaces in the MAX TNT. Use of this capability is controlled via IPDC messages from the signaling gateway.

SS7 link establishment timer

With MAX TNT TAOS 8.0.0, the MAX TNT supports a T5 timer that automatically enables link connection and reconnection requests to the signaling gateway to occur at random intervals. The timer can help prevent the signaling gateway from receiving many link connect

requests within a short period of time, especially when the signaling gateway is connected with many MAX TNT units.

After its link to the gateway is disconnected, the MAX TNT initializes the T5 timer with a random value between 0 and 6 seconds and attempts a connection when the timer expires. After each failed connection attempt, the MAX TNT increases the T5 time-out value by 1 second until it reaches 20 seconds. The timer remains at 20 seconds for subsequent connection attempts.

The MAX TNT resets the T5 timer as soon as the link is active.

Support for 2-wire continuity check on T1 lines

Initially, MAX TNT units supported a 4-wire-only continuity check as defined in Q.724 Sections 7 and 8, ANSI T1.113.4 Annex B, GR-246-CORE Annex B. The 4-wire continuity check requires one end of a line to place a channel into loopback state while the other end sends a tone. The check concludes successfully if the tone sent on the outgoing path is received on the return path within acceptable transmission and timing limits. The 4-wire check procedure cannot detect potential inadvertent loops in the line path or in line facilities, and cannot be used when the other exchange is analog. For these reasons, the procedure known as 2-wire continuity check is recommended by the International Telecommunications Union Telecommunication Standardization Sector (ITU-T).

Note: The ITU-T carries out the operations of the former Consultative Committee for International Telephone and Telegraph (CCITT).

With MAX TNT TAOS 8.0.0, the MAX TNT supports both incoming and outgoing 2-wire continuity checks for T1 lines only. You can select the type of check to perform on a per-line basis. Both the native 2-wire continuity check (GR-246-CORE Section B.2) and 4-wire-to-2-wire emulation (GR-246-CORE Section B.3) are supported.

Outgoing continuity tests are supported only on T1 and T3 cards. E1 cards support receipt of 4-wire continuity check requests only, and cannot originate continuity tests.

An SS7-Continuity subprofile has been added to the T1 profile to allow you to specify the type of incoming and outgoing continuity checks to perform for all channels on a line. Both ends of the connection must agree on the continuity check to be used for the line. Following are the relevant parameters, shown with default values:

```
[in T1/{ shelf-1 slot-1 1 }:line-interface:ss7-continuity]
incoming-procedure = loopback
outgoing-procedure = single-tone-2010
```

Parameter	Specifies
Incoming-Procedure	Loopback or transponder test mode. The <code>loopback</code> setting (the default) places the channel into loopback mode during the continuity test. This mode must be used if the line is provisioned for an incoming 4-wire continuity test. The <code>transponder</code> setting places the channel into Tone Transponder mode during the continuity test. In this mode, the channel can detect two tones: 2010Hz and 1780Hz. When either tone is detected, the other one is returned. This mode should be used for lines provisioned for incoming 2-wire and 4-wire-to-2-wire continuity checks.

Parameter	Specifies
Outgoing-Procedure	<p>Type of continuity check. With the <code>single-tone-2010</code> setting (the default), the MAX TNT sends a 2010Hz tone and expects to receive a 2010Hz tone in return. This procedure is generally known as a 4-wire continuity check.</p> <p>With the <code>send-2010-expect-1780</code> setting, the MAX TNT sends a 2010Hz tone and expects to receive 1780Hz tone in return. This procedure is generally known as a 2-wire continuity check.</p> <p>With the <code>send-1780-expect-2010</code> setting, the MAX TNT sends a 1780Hz tone and expects to receive a 2010Hz tone in return. This procedure is generally known as a 4-wire to 2-wire continuity check.</p> <p>If you change the type of a continuity check, the new type is used for new continuity check requests on the line as soon as the line profile is saved. Existing check-loops that are already active on the line are not modified or canceled when the profile is saved.</p>

The type of the continuity check procedure to be used is determined by line provisioning and is agreed upon by the connecting exchanges. SS7 signaling procedures used for continuity check (Q.764 Section G.3, ANSI T1.113.4 Section 2.1.6) are the same for both 4-wire and 2-wire circuits, but the behavior of trunk termination devices is different.

The native 2-wire continuity check procedure requires that the loopback be replaced by a transponder and that a 1780Hz \pm 20Hz tone be used in the return direction.

The MAX TNT also supports the 4-wire-to-2-wire continuity check, with the following requirements: The exchange that terminates 4 wires must use a transmitting frequency of 1780 \pm 20Hz and a receiving frequency 2010 \pm 30Hz. The exchange that terminates the 2 wires must use a transmitting frequency of 2010 \pm 8Hz and a receiving frequency of 1780 \pm 30Hz.

Outgoing continuity tests on T1 and T3

Initially, MAX TNT units supported incoming continuity tests only. During these tests, the telephone switch requests that the MAX TNT put a DS0 channel into a loopback and then generates a 2010Hz tone. If the switch receives the tone in return, the continuity test is successful.

With MAX TNT TAOS 8.0.0, the MAX TNT also supports outgoing call continuity tests on T1 and T3 cards. For outgoing continuity, the switch puts a DS0 into a loopback and the MAX TNT generates a 2010Hz tone. If the MAX TNT receives the tone in return, the continuity test is successful. Note that all the setup and signaling required to coordinate a continuity test is handled by the signaling gateway via SS7.

Note: Outgoing continuity tests are supported only on T1 and T3 cards. E1 cards support receipt of 4-wire continuity check requests only, and cannot originate continuity tests.

Digital milliwatt tone support on T1 and T3

With MAX TNT TAOS 8.0.0, T1 and T3 cards generate the 1000Hz digital milliwatt (DMW) tone. The SS7 switch sends a digital milliwatt tone request to the MAX TNT over IPDC and

uses the tone that the MAX TNT generates in special test calls to measure the line distortion and attenuation in the telephone network.

Analog milliwatt tone and variable tone support

With MAX TNT TAOS 8.0.0, changes were made to IPDC Tone-Type and Tone-String tags to enable the IPDC Specify Tone (STN) message to generate the analog milliwatt tones. When the MAX TNT receives a message from the signaling gateway specifying the new tags, it responds with the appropriate tone type or tone string.

When the signaling gateway specifies a variable tone, it details the tone in the IPDC Tone-String message tag, which uses the following format:

"frequency1, frequency2, amplitude, duration"

Element	Description
Frequency1	First frequency of the dual tone. This value can range from 1 to 3999 and has an accuracy of ± 1 Hz.
Frequency2	The second frequency of the dual tone. This value can range from 0 for single tones to 3999 and has an accuracy of ± 1 Hz.
Amplitude	Amplitude of the tone. If this value is in the range of 4 through 32767, it is an absolute value. A value in the range from -49 through 2 is a decibel level. The following relationship exists between decibel levels and absolute values: $\text{dBm0} = 20 * \log_{10} (\text{absolute value} / 22748.4)$
Duration	Duration of the tone in milliseconds. This value can range from 0 to 2631. If the duration is 0, a tone will be played continuously until it is stopped by a second STN command.

For example, the following string defines a 1004hz tone at 22748 amplitude for 1 second:

"1004, 0, 22748, 1000"

The following string defines a dual tone with frequency 697Hz and 1477Hz at 14567 amplitude for 2 seconds:

"697, 1477, 14567, 2000"

The following string defines a 2050, -3dBm0 tone played continuously until stopped by a second STN message:

"2050, 0, -3, 0"

IPDC enhancements for reporting VoIP call statistics

With MAX TNT TAOS 8.0.0, a MAX TNT operating as a network access server (NAS) with a signaling gateway can report VoIP call statistics in the output of the NAS messaging interface. IPDC VoIP call statistics are reported once a call is cleared. The source that originates call clearing can be either the signaling gateway or the MAX TNT.

Supported tags for reporting statistics

IPDC 0.12 statistics tags are reported when the signaling gateway or the MAX TNT clears calls under the following conditions:

- When the access server initiates a call teardown using an RCR message.
- For packet-based calls when the access server acknowledges a call teardown using an ACR message

The MAX TNT reports the following VoIP statistics, as defined by IPDC 0.12:

- Number of Real-Time Protocol (RTP) audio packets sent and received by the MAX TNT.
- Number of RTP audio packets that failed to reach the MAX TNT as determined by missed sequence numbers.
- Number of audio bytes in the RTP payload sent by the MAX TNT.
- Number of audio bytes received in the RTP payload that failed to reach the MAX TNT. Because the number of bytes per packet is variable, this value can only be estimated, based upon an average packet size multiplied by the number of nonreceived packets. This value can also be estimated by the control server with the information supplied.
- Number of RTP audio packets received.
- Number of audio bytes received in the RTP payload.
- Estimated interarrival jitter (in milliseconds) Interarrival jitter is an estimate of the statistical variance among the arrival times of RTP packets, which is equivalent to the difference in their relative transit times. Relative transit time is the difference between a packet's RTP timestamp at the sender and the receiver's clock at the time of arrival.

ss7nmi debug-level command

The MAX TNT reports the VoIP call statistics in the output of the `ss7nmi debug-level` command. When the command is entered with the `-s` option, the results displayed include the number of release channel request (RCR) and release channel completed (ACR) messages sent with and without VoIP call statistics, and the number of unknown SS7 VoIP messages. In the following example, new statistics reported for IPDC VoIP calls are shown in bold type:

```
admin> ss7nmi -s
SS7 NAS Messaging Interface (NMI) statistics:
    Initialized successfully:                Yes
    Total number of internal errors:         0
    Level of diagnostics:                   0
Signaling Layer:
    Current link state:                     STARTING
    Last generated transaction ID:           1
    Timer T305 (RST1):                     1000 ticks - idle
    Number of protocol version errors:       0
    Number of 'message reject' received:     0
    Number of bad packets received:          0
    Number of unknown messages:             0
    Number of unknown SS7Voip messages:    0
    Number of resource conflicts:            0
    Number of release race conditions:       0
    Number of RCR with stats sent:         0
    Number of RCR without stats sent:      0
    Number of ACR with stats sent:        0
    Number of ACR without stats sent:     0
Data Transport Layer:
    Number of link fail-overs:              0
    Number of persistent errors:            0
    Last error:                             No Error
    Last error timestamp:                   [01/01/1990 00:00:00]
```

Statistics and error reporting on SS7 connections

The `ss7asg -s` command provides detailed interface information about statistics and error conditions on SS7 connections. The output differs depending on whether errors are detected.

Note: The `ss7asg -r` command resets all the signaling layer statistics to 0 and updates the timestamp to the time the counters were reset.

Command output when no errors are detected

The following sample output indicates that no errors were detected in SS7 connections:

```
admin> ss7asg -s
SS7 Signaling Gateway interface statistics:
    Initialized successfully:          Yes
    Interface state:                  Enabled/Down
    Diagnostic level:                  0

Signaling Layer:
    Number of SETUP requests from:    L2: 0          CC: 0
    Number of CONNECT to ASG:         0
    Number of CONNECT_ACK from ASG:    0
    Number of SETUP rejected from:     L3: 0          CC: 0
    Number of DISCONNECT requests from: L2: 0          CC: 0
    Number of REGISTRATION to ASG:     0
    Number of REGISTRATION_ACK from ASG: 0
    Number of DL_REL_IND from L2:      0
    Number of DL_EST_IND from L2:      0
    Number of T303 expiry events:      0
    Number of T305 expiry events:      0
    Number of T308 expiry events:      0
    Last L3 counters reset timestamp:   [02/08/1999 18:47:41]

Data Transport Layer:
    Number of link fail-overs:         0
    Number of persistent errors:       161
    Last error:                        Persistent Error
    Last error status change timestamp: [02/08/1999 18:47:41]
```

When the command reports no errors, the output contains the following fields:

Output field	Description
Initialized successfully	Indicates whether the SS7 layer between the MAX TNT and the signaling gateway has been successfully initialized.
Interface state	State of the SS7 interface. A value of Enabled/Up indicates that the Enabled parameter in the SS7-Gateway profile is set to yes. A value of Enabled/Down indicates that the Enabled parameter in the SS7-Gateway profile is set to yes, but the TCP link to the signaling gateway is down. A value of Disabled indicates that the Enabled parameter in the SS7-Gateway profile is set to no.

Output field	Description
Diagnostic level	<p>The diagnostic level as specified with the <code>-t</code> option. Values can be one of the following:</p> <ul style="list-style-type: none"> • 0: Disable diagnostic output. • 1: Show errors only. • 2: Trace L3 events and states. • 3: Trace Call Control events. • 4: Show all task events. • 5: Dump L3 packets. • 6: Dump Call Control primitives.
Number of SETUP requests from:	<p>L2: Number of setup requests from the signaling gateway (SS7 network) or from incoming calls.</p> <p>CC: Number of times the MAX TNT tried to make an outgoing call to the signaling gateway (the SS7 network). Note that outgoing calls are not currently supported.</p>
Number of CONNECT to ASG	Total number of active connections to the signaling gateway since it was last reset.
Number of CONNECT_ACK from ASG	Number of connection acknowledgements the MAX TNT has received from the signaling gateway.
Number of SETUP rejected from:	<p>Number of setup requests rejected by layer 3 and the signaling gateway call control.</p> <ul style="list-style-type: none"> • Setups rejected by L3 indicate a packet decode error on the incoming setup request. • Setups rejected by CC can mean that no route or resource exists, or that authentication failed for the incoming call.
Number of DISCONNECT requests from:	<p>Number of disconnection requests from layer 2 and the signaling gateway call control.</p> <ul style="list-style-type: none"> • Disconnection requests from layer 2 are initiated by the signaling gateway. • Disconnection requests from CC are initiated by the MAX TNT.
Number of REGISTRATION to ASG	Number of registration requests the MAX TNT has sent to the signaling gateway.
Number of REGISTRATION_ACK from ASG	Number of registration acknowledgments the MAX TNT has received from the signaling gateway.
Number of DL_REL_IND from L2	Number of Data Link Release Indication messages received from layer 2. Layer 2 sends these messages to layer 3 to inform it about the status of the link. Data Link Release Indication messages mean that the link between the MAX TNT and the signaling gateway is down and communication is not possible.

Output field	Description
Number of DL_EST_IND from L2	Number of Data Link Establish Indication messages received from layer 2. Layer 2 sends these messages to layer 3 to inform it about the status of the link. Data Link Establish Indication messages mean that the link between the MAX TNT and the signaling gateway has been reestablished and communication is possible.
Number of T303 expiry events	Number of times the T303 timer expired.
Number of T305 expiry events	Number of times the T305 timer expired.
Number of T308 expiry events	Number of times the T308 timer expired.
Last L3 counters reset timestamp	Time the signaling layer timers were last reset using the <code>ss7asg -r</code> command.
Number of link fail-overs	In a dual LAN configuration, the number of times the MAX TNT switched from one TCP/IP messaging link to another due to the failure of the link.
Number of persistent errors	Number of times the MAX TNT tried to reestablish a layer 2 link.
Last error	Type of last error. Possible values are: <ul style="list-style-type: none"> • No Error: L2 is operating normally. • Link Loss: Link down. • Persistent Error: Link down. • Link Shutdown: Link disabled. • Link Fail-over: Switched to secondary LAN connection.
Last error status change timestamp	Time the last error occurred.

Command output showing errors

The following sample indicates that errors were detected in SS7 connections:

```
admin> ss7asg -s
SS7 Signaling Gateway interface statistics:
    Initialized successfully:      Yes
    Interface state:              Enabled/Down
    Diagnostic level:              0

Errors:
    Number of memory allocation failures:  0
    Number of errors in profile operations:  0
    Number of invalid memory pointers:      0
    Number of internal errors:              8

Initialization Errors:
    Number of errors in initialization:      8
    Memory pools:                           0
    Mailboxes:                              0
```

Extension features in MAX TNT TAOS 8.0.0

Signaling System 7 (SS7)

```
Signaling Layer:
  Number of SETUP requests from:           L2: 0          CC: 0
  Number of CONNECT to ASG:                0
  Number of CONNECT_ACK from ASG:          0
  Number of SETUP rejected from:           L3: 0          CC: 0
  Number of DISCONNECT requests from:      L2: 0          CC: 0
  Number of REGISTRATION to ASG:           0
  Number of REGISTRATION_ACK from ASG:     0
  Number of DL_REL_IND from L2:             0
  Number of DL_EST_IND from L2:             0
  Number of T303 expiry events:             0
  Number of T305 expiry events:             0
  Number of T308 expiry events:             0
  Last L3 counters reset timestamp:         [02/16/1999 10:33:31]

Data Transport Layer:
  Number of link fail-overs:                0
  Number of persistent errors:              0
  Last error:                              No Error
  Last error status change timestamp:       [01/01/1990 00:00:00]
```

When errors are detected, the command output displays the fields explained in the previous section plus the following additional information:

Output field	Description
Number of memory allocation failures	Number of times the MAX TNT could not allocate memory for packets traveling between call control and layer 3. These errors might occur if the MAX TNT does not have a 32-MB DRAM card installed.
Number of errors in profile operations	Number of times the MAX TNT could not register the SS7-Gateway profile or read or update a T1 profile.
Number of invalid memory pointers	Number of empty packets received by IPDC layer 3. Used for IPDC only.
Number of internal errors	Number of internal errors.
Number of errors in initialization	Number of errors that occurred during the initialization of the SS7 ASG interface.
Memory pools	Number of buffer pool allocations that failed.
Mailboxes	Number of failures that occurred during the creation or operation of the mailboxes used for interlayer messaging.

Cause codes for SS7 ASGCP calls to the MAX TNT

With MAX TNT TAOS 8.0.0, the MAX TNT reports cause codes to the signaling gateway via ASGCP when it initiates a call clearing. The following ASGCP messages carry cause code information.

- Disconnect
- Release
- Release Complete
- Restart Acknowledgement (cause optional)

- Status

The MAX TNT currently reports the cause codes defined by ITU-T Recommendation Q.850. For definitions of the individual cause values, refer to Q.850. Note the following:

- The MAX TNT reports Normal call clearing (16) if a MAX TNT modem times out on a modem call.
- The MAX TNT reports User busy (17) if the MAX TNT cannot find a route, or if no resource is available for the call.

SS7 IPDC support for call ID and disconnect cause codes

With MAX TNT TAOS 8.0.0, the MAX TNT reports a globally unique call identifier to call-logging servers for SS7 data or VoIP calls. This feature enables the NavisAccess software to associate call statistics information generated by the signaling gateway and by the MAX TNT.

A similar mechanism is supported in the H.323 VoIP context, where a well-defined globally unique call ID is set by the originating endpoint. This call ID is used to associate RAS signaling with the modified Q.931 call control signaling used in H.225.0 call setup. In an H.323 VoIP environment, the MAX TNT reports the call ID to call-logging servers when a call is connected, maintained, and terminated. (H.323, Q.931, and H.225.0 are ITU-T recommendations for voice communication over networks.) For more information about H.323 VoIP see “Overview of VoIP in an H.323v2 environment” on page 257.

To support this functionality in the SS7 IPDC context, the following changes were made:

- IPDC now generates a globally unique call ID for SS7 VoIP and data calls.
- IPDC now includes the globally unique call ID in IPDC messages.
- The MAX TNT now reports the call ID to call-logging servers.

IPDC generation of a globally unique call ID

IPDC uses the same definition and algorithm for generating a globally unique callIdentifier as H.225.0. The ID consists of a record of 16 octets. For details, refer to H.225.0, Version 2, pages 44 to 47. The signaling gateway creates the call ID in the following cases:

- Request inbound call setup (RCSI) message
- Request pass-through call setup for TDM connection between two channels (RCST) message
- Request packet pass-through call (RCCP) message

The MAX TNT creates a call ID for a request outbound call setup (RCSO) message. Note that the MAX TNT does not currently report the call ID of outbound calls to call-logging servers.

New Global-Call-ID parameter

A new parameter has been added to the Call-Info profile to report the global call ID. The Global-Call-ID parameter is shown with a sample setting in the following example:

```
[in CALL-INFO/{ 3 }]  
mbid* = { 3 }  
call-service = switched  
called-number-type = 2  
nailed-up-group = 1
```

```
call-by-call = 0
phone-number = ""
transit-number = ""
billing-number = ""
switched-call-type = 67
ftl-caller = 0
calling-number = { "" unknown unknown unspecified unspecified }
force-56kbps = 0
redirect-number = ""
call-direction = 0
global-call-id = 03040506-0102-0900-0807-010203040506
```

Modification of Start and Stop records

The Ascend_Global_Call_Id attribute has been added to Start and Stop records for SS7 VoIP and data calls. The attribute is added for call-logging only, not RADIUS, and is reported only when the global call ID is available.

The MAX TNT sends Stop records for SS7 calls that are cleared or rejected at the SS7 IPDC layer. Those calls do not have Start records, because they are never routed to host cards.

New disconnect cause codes

The following set of disconnect cause codes reports the cause of termination for calls that are cleared or rejected at the SS7 IPDC layer. These codes are based on the cause codes defined by ITU-T Recommendation Q.850, *Usage of Cause and Location in the Digital Subscriber Signaling System No. 1 and the Signaling System No. 7 ISDN User Part*. This group of cause codes begins at offset 800.

Event	Code	Q.850 Definition
DIS_Q850_UNASSIGNED_NUMBER	801	Unallocated (unassigned) number
DIS_Q850_NO_ROUTE	802	No route to specified transit network
DIS_Q850_NO_ROUTE_TO_DEST	803	No route to destination
DIS_Q850_CHANNEL_UNACCEPTABLE	806	Channel unacceptable
DIS_Q850_NORMAL_CLEARING	816	Normal call clearing
DIS_Q850_USER_BUSY	817	User busy
DIS_Q850_NO_USER_RESPONDING	818	No user responding
DIS_Q850_USER_ALERT_NO_ANSWER	819	No answer from user (user alerted)
DIS_Q850_CALL_REJECTED	821	Call rejected
DIS_Q850_NUMBER_CHANGED	822	Number changed
DIS_Q850_DEST_OUT_OF_ORDER	827	Destination out of order
DIS_Q850_INVALID_NUMBER_FORMAT	828	Invalid number format (address incomplete)
DIS_Q850_FACILITY_REJECTED	829	Facility rejected
DIS_Q850_RESP_TO_STAT_ENQ	830	Response to STATUS ENQUIRY
DIS_Q850_UNSPECIFIED_CAUSE	831	Unspecified normal event
DIS_Q850_NO_CIRCUIT_AVAILABLE	834	No circuit or channel available
DIS_Q850_NETWORK_OUT_OF_ORDER	838	Network out of order

Event	Code	Q.850 Definition
DIS_Q850_TEMPORARY_FAILURE	841	Temporary failure
DIS_Q850_NETWORK_CONGESTION	842	Switching equipment congestion
DIS_Q850_ACCESS_INFO_DISCARDED	843	Access information discarded
DIS_Q850_REQ_CHANNEL_NOT_AVAIL	844	Requested circuit or channel not available
DIS_Q850_PRE_EMPTED	845	Call preempted
DIS_Q850_RESOURCE_NOT_AVAIL	847	Resource unavailable
DIS_Q850_FACILITY_NOT_SUBSCRIBED	850	Requested facility not subscribed
DIS_Q850_OUTGOING_CALL_BARRED	852	Outgoing calls barred within the CUG
DIS_Q850_INCOMING_CALL_BARRED	854	Incoming calls barred within the CUG
DIS_Q850_BEAR_CAP_NOT_AVAIL	858	Bearer capability not presently available
DIS_Q850_SERVICE_NOT_AVAIL	863	Service or option not available, unspecified
DIS_Q850_CAP_NOT_IMPLEMENTED	865	Bearer capability not implemented
DIS_Q850_CHAN_NOT_IMPLEMENTED	866	Channel type not implemented
DIS_Q850_FACILITY_NOT_IMPLEMENT	869	Requested facility not implemented
DIS_Q850_INVALID_CALL_REF	881	Invalid call reference value
DIS_Q850_CHAN_DOES_NOT_EXIST	882	Identified channel does not exist
DIS_Q850_INCOMPATIBLE_DEST	888	Incompatible destination
DIS_Q850_MANDATORY_IE_MISSING	896	Mandatory information element missing
DIS_Q850_NONEXISTENT_MSG	897	Message type nonexistent or not implemented
DIS_Q850_WRONG_MESSAGE	898	Message not compatible with call state, or message type nonexistent or not implemented
DIS_Q850_NONEXISTENT_IE	899	Information element or parameter nonexistent or not implemented
DIS_Q850_INVALID_ELEM_CONTENTS	900	Invalid information element contents
DIS_Q850_WRONG_MSG_FOR_STAT	901	Message not compatible with call state
DIS_Q850_TIMER_EXPIRY	902	Recovery on timer expiration
DIS_Q850_MANDATORY_IE_LEN_ERR	903	Parameter that was nonexistent or not implemented was passed on
DIS_Q850_PROTOCOL_ERROR	911	Message with unrecognized parameter was discarded
DIS_Q850_INTERWORKING_UNSPEC	927	Unspecified internetworking event

SNMP: Support for the SS7 MIB (ascend 29)

In MAX TNT TAOS 8.0.0, a new SS7 MIB is supported. The MIB file is `mgstat.mib`. The MIB is implemented as a branch object with the main object, `mgGroup`, linked into the Ascend enterprise MIB. For definitions and descriptions of objects, see the `mgstat.mib` file distributed with MAX TNT TAOS 8.0.0.

SNMP: Support for SS7 link-state trap

In MAX TNT TAOS 8.0.0, an SNMP trap is supported for reporting the status of the link between SS7 media gateways and the MAX TNT. The trap can be configured when an SS7 license is enabled. For a trap to be generated when the trap condition occurs, SNMP traps must be enabled and the setting for the trap condition must be enabled. For details about enabling traps, see the *MAX TNT Administration Guide*.

The following trap has been added to the Ascend enterprise traps:

```
megacoLinkStatusTrap  TRAP-TYPE
    ENTERPRISE          ascend
    VARIABLES            { mgLinkName, mgOperStatus }
    DESCRIPTION          "This trap indicates that operational status
                        of a media gateway control link has changed."
    ::= 42
```

Following is the relevant parameter in the Trap profile, shown with its default value, for enabling the trap:

```
[in TRAP/""]
megaco-link-status-enabled = no
```

Parameter	Specifies
Megaco-Link-Status-Enabled	Enable/disable trap generation of communication link status between the SS7 media gateway and the MAX TNT. This trap indicates that operational status of a media gateway control link has changed from any state to the Up state or from Up state to any other state. Changes to this parameter become effective when you write the Trap profile. The trap contains the name of the link, which is currently always reported as <code>default</code> , and the new operational status.

For example, the following commands enable the SS7 link-state trap:

```
admin> read trap example
TRAP/example read

admin> set megaco-link-status-enabled = yes

admin> write
TRAP/example written
```

SNMP: Idle time attribute in active session table

In MAX TNT TAOS 8.0.0, a new `ssnActiveIdleTime` attribute has been added to the `sessionActiveTable` to show the number of ticks (0.01-second intervals) during which the current session has been idle. The new attribute has the object ID

`sessionActiveEntry.8`. No other object ID was modified in the `sessionActiveTable`.

MultiVoice operations

Support for MultiVoice operations was introduced with limited availability in earlier TAOS 7.x releases. It is now generally available in MAX TNT TAOS 8.0.0.

MultiVoice functionality includes Voice over IP (VoIP) and a transparent data mode that enables users to run a modem on a VoIP channel. With a separate license on both ends of the transmission, MultiVoice also supports real-time fax over IP, which conforms to the T.38 ITU-T standard.

Note: This release note provides an overview of MultiVoice functionality and describes new MAX TNT TAOS 8.0.0 features that are not documented in the *MultiVoice for the MAX TNT Configuration Guide*. For details about MultiVoice configuration, see the guides at <http://www.ascend.com/doclibrary>.

In MAX TNT TAOS 8.0.0, the following MultiVoice software licenses can be enabled:

- VoIP, which enables the MAX TNT to act as an H.323v2 MultiVoice Gateway for transmission of real-time voice calls and transparent modem calls across IP networks.
- VoIP and SS7, which enables the MAX TNT to act as a MultiVoice Gateway that communicates with an SS7 signaling gateway to transmit real-time voice calls and transparent modem calls from an SS7 network across IP networks.
- Real-time fax (T.38) over IP, which uses the VoIP framework for call establishment, fax detection, and fax initiation.

For information about using MultiVoice for basic long-distance service and 800 service, and with overlapping coverage areas and multizone call routing, see the *MultiVoice for the MAX TNT Configuration Guide* at <http://www.ascend.com/doclibrary>.

System requirements for VoIP

To operate as a MultiVoice Gateway, a MAX TNT unit must have the following equipment and licenses:

- VoIP software licenses
- Sufficient MultiDSP cards to process VoIP calls
- Sufficient T1, T3, or E1 trunks to process VoIP calls
- Sufficient Ethernet-3 cards to process VoIP calls

If the MAX TNT unit will operate in an H.323 environment, it must also have an IP connection to a workstation running the MultiVoice Access Manager (MVAM) software.

If the unit will operate in an SS7 environment, the SS7 software license must also be enabled so that the system can perform IPDC packet processing. For details about SS7 requirements, see “Signaling System 7 (SS7)” on page 233.

Ethernet requirements for VoIP processing

MAX TNT units do not support routing of VoIP calls through the shelf controller Ethernet port. Ethernet-3 (TNT-SL-E100-V-C) cards are required for VoIP. The Ethernet-3 card is a high performance Ethernet module with one 100-MB interface designed for demanding applications such as VoIP.

Full-duplex mode required

When using the Ethernet-3 card to support VoIP call processing, the card must operate in full-duplex mode. The card operates in full-duplex mode by default, as specified in the setting of the following parameter:

```
[in ETHERNET/{ any-shelf any-slot 0 }]
duplex-mode = full-duplex
```

Compatible configuration in connecting port of hub or router

The 100-MB interface on the Ethernet-3 card is not autoconfigurable and Lucent does not recommend connecting it to a hub or router port that has been autoconfigured. Connecting it to an autoconfigured port can have negative effects on VoIP calls, including poor voice quality for connected calls and increased instances of initial call failures.

To ensure the best performance and quality for VoIP calls, make sure that the hub or router port that connects the Ethernet-3 card to the packet network complies with the following recommended configuration:

- Port autoconfiguration must be disabled.
- Port speed must be configured to operate at 100 Mbits only.
- Port must be configured for full-duplex transmission

Please refer to the manufacturer-provided documentation for your particular network hub, router or switch for specific instructions on configuring its Ethernet ports.

Note: It is not necessary to apply the recommended configuration to ports providing the outbound connection from the hub or router to the rest of the IP network. This configuration is required only for the port connecting to the Ethernet-3 card.

VoIP call routing

When a VoIP license has been enabled, the system creates a new Call-Route profile for each installed MultiDSP card that supports VoIP. The new Call-Route profile sets the Call-Route-Type parameter to `voip-call-type`, as shown in the following sample profile for a MultiDSP card in shelf 1, slot 3:


```
admin> get call-route { { { 1 3 0 } 0 } 2 }
[in CALL-ROUTE/{ { { shelf-1 slot-3 0 } 0 2 } ]
index* = { { { shelf-1 slot-3 0 } 0 } 2 }
trunk-group = 0
phone-number = ""
preferred-source = { { any-shelf any-slot 0 } 0 }
call-route-type = voip-call-type
```

The `voip-call-type` setting enables the system to route VoIP calls to the MultiDSP card. When the MAX TNT receives a VoIP call on a network line (such as T1 or E1), it routes the

traffic internally on its time-division multiplex (TDM) bus to the MultiDSP card, which handles VoIP-related functions such as audio coder/decoder (codec) processing, RTP and UDP processing, and so forth.

The MultiDSP card then forwards the packetized traffic on the system's packet bus to an exit (egress) interface such as Ethernet or another T1 line. The example path shown in Figure 40 provides a simplified picture of how VoIP calls are routed through the MAX TNT.

Figure 40. Simplified view of VoIP call routing within the MAX TNT



For details about VoIP call routing and how to fine tune it, see the *MultiVoice for the MAX TNT Configuration Guide* at <http://www.ascend.com/doclibrary>.


Overview of VoIP in an H.323v2 environment

MultiVoice is compliant with the ITU-T H.323 standard for the transmission of real-time voice communications across IP networks. H.323 systems use the IETF standard Real-Time Transport Protocol (RTP) with codecs for voice and other communications over the Internet.

VoIP-enabled MAX TNT units operate as MultiVoice Gateways. Callers dial into a local MAX TNT through the PSTN. The MAX TNT then communicates with a MultiVoice Gatekeeper to establish communication channels to a far end MultiVoice Gateway. Workstations running MVAM software operate as H.323 MultiVoice Gatekeepers, which handle all call control functions, including bandwidth control, authentication, call-detail recording (CDR), and alias translation.

In the example Gateway and Gatekeeper configuration in Figure 41, two Gateways connect Caller A to Caller B. A system running MVAM performs the H.323 Gatekeeper functions.

Figure 41. Example diagram of MultiVoice in H.323 environment



When Caller A dials Caller B, events such as the following occur:

- 1 Caller A dials Gateway 1, and enters his or her PIN authentication (if required) and Caller B's telephone number.
- 2 Gateway 1 establishes a session with the Gatekeeper, and then forwards the telephone number and PIN authentication to the Gatekeeper.
- 3 The Gatekeeper authenticates Caller A and, if authentication is successful, forwards the IP address of Gateway 2 to Gateway 1.
- 4 Gateway 1 establishes a session with Gateway 2.
- 5 Gateway 2 forwards the call request to Caller B.
- 6 When Caller B answers the telephone (goes off-hook), voice traffic is tunneled in IP packets by means of RTP, between Gateway 1 and Gateway 2.


Overview of VoIP in an SS7 IPDC 0.12 environment

In an SS7 environment, VoIP-enabled MAX TNT units are MultiVoice Gateways that communicate with an SS7 signaling gateway to establish communication channels to a far-end MultiVoice Gateway.

The SS7 signaling gateways initiate and manage call setup and release, and execute call routing. The signaling gateway communicates call setup information to the MAX TNT using IPDC 0.12. IPDC message tags define voice encoding type, packet loading, IP and RTP ports, and other variables used for processing VoIP calls

In the example MultiVoice Gateway and signaling gateway configuration in Figure 42, the Gateways support VoIP calls controlled by IPDC over intermachine trunks (IMTs) for SS7 calls originating from the PSTN.

Figure 42. Example diagram of MultiVoice in SS7 environment



When Caller A dials Caller B, the following events occur:

- 1 Caller A dials the number for their SS7 service provider plus Caller B's telephone number. For example, Caller A dials a number such as 10-10-999-1-888-555-1212.
- 2 The signaling gateway assembles call routing information, and other information required to connect the call, such as user authentication and call reporting information.
- 3 The signaling gateway then sends an SS7 message to the PSTN to ring Caller B's telephone.
- 4 The signaling gateway uses IPDC to initiate an RTP/IP connection across the packet network between Gateway 1 and Gateway 2. The signaling gateway simultaneously sends IPDC setup information to both Gateway 1 and Gateway 2.
- 5 When Caller B answers the telephone (goes off-hook), the signaling gateway converts the SS7 signals into IPDC packets, and voice traffic is tunneled in IP packets between Gateway 1 and Gateway 2 by means of RTP.
- 6 Gateway 2 passes the IPDC packets to the signaling gateway at the far end, which converts the IPDC packets to SS7 messages and routes the call across the appropriate signaling links to Caller B.

In an SS7 environment, values in IPDC message tags override corresponding call management settings in the default VoIP profile. For details about configuring a MAX TNT unit to interact with an SS7 signaling gateway using IPDC 0.12, see "Signaling System 7 (SS7)" on page 233.

General system configuration for VoIP support

Lucent recommends certain IP and call-handling configurations for processing VoIP calls. Global settings that are required for VoIP communication are also described in this section.

Note: For details about recommended IP settings and routes, see the *MultiVoice for the MAX TNT Configuration Guide* at <http://www.ascend.com/doclibrary>.

Disabling ICMP Destination Unreachable packets for VoIP calls

For Voice over IP (VoIP) calls, UDP for-me packets can arrive at a rate of 200 packets per second for each direction of each call. If the MAX TNT is not listening on a port for the for-me

packets while setting up or tearing down a call, it returns ICMP Destination Unreachable packets at the same rate. To prevent the performance penalty caused by this situation, you can now configure the system not to send ICMP Destination Unreachable packets.

Caution: This feature is intended only for VoIP environments. Enabling this feature can break required behavior for IPv4 routers, such as Path MTU Discovery.

Following is the relevant parameter, shown with its default setting:

```
[in IP-GLOBAL]
send-icmp-dest-unreachable = yes
```

Parameter	Specifies
Send-ICMP-Dest-Unreachable	Enable/disable sending of ICMP Destination Unreachable packets. The default is <i>yes</i> . If set to <i>no</i> , the MAX TNT does not send ICMP Destination Unreachable packets. Setting this parameter to <i>No</i> is recommended only for VoIP environments.

The following commands disable transmission of ICMP Destination Unreachable packets:

```
admin> read ip-global
IP-GLOBAL read

admin> set send-icmp-dest-unreachable = no

admin> write
IP-GLOBAL written
```

Preventing receipt of UDP packets until VoIP calls are set up

When two MultiVoice Gateway systems are establishing the link for transmission of a VoIP call, both systems do not always complete the call setup at the same time. However, a Gateway starts sending UDP packets to the other Gateway as soon its own call setup is complete. If the receiving Gateway has not yet set up its port caches, the shelf controller receives the UDP packets for a period of time until the call is fully set up. Now, you can prevent receipt of UDP packets until the link is fully established. Following is the relevant parameter, shown with the default value:

```
[in IP-GLOBAL]
throttle-no-port-match-udp-traffic-on-slot = no
```

Parameter	Specifies
Throttle-No-Port-Match-UDP-Traffic-On-Slot	Enable/disable reception of UDP packets for UDP ports currently unknown to the MAX TNT. With the default value of <i>no</i> , the system behaves as in previous releases and sends the unknown port packets to the shelf controller for processing. If the parameter is set to <i>yes</i> , the system discards UDP packets until the UDP port is known. The setting of <i>yes</i> is recommended for MultiVoice Gateways, to prevent overloading of the shelf controller when both Gateways do not always complete the VoIP call setup at the same time.

The following commands enable the system to discard UDP packets until the UDP port is known:

```
admin> read ip-global
IP-GLOBAL read

admin> set throttle-no-port-match-udp-traffic-on-slot = yes

admin> write
IP-GLOBAL written
```

System settings for VoIP operations

Lucent recommends setting the following parameters, shown with default values, to facilitate VoIP call handling:

```
[in IP-GLOBAL]
system-ip-addr = 0.0.0.0

[in ANSWER-DEFAULTS:session-info]
idle-timer = 0

[in SYSTEM]
max-dialout-time = 60
parallel-dialing = 32
country = us
```

Parameter	Recommended VoIP settings
System-IP-Addr	In an H.323 environment, set this parameter to the shelf controller IP address. In an IPDC environment, if the system allocates its own listen address, set this parameter to the IP address of a LAN interface other than the shelf controller port.
Idle-Timer	For real-time fax or transparent modem calls, set this parameter should be set to 0 to disable the idle timer and prevent the fax or modem calls from timing out.
Max-Dialout-Time	To allow sufficient time for the MAX TNT to establish the connection to the called destination, and for consistency with internal H.323 timers, a setting of 60 is recommended.
Parallel-Dialing	To decrease the instances when VoIP callers wait for a silent interval while the MAX TNT completes a call that has been queued, a setting of 32 is recommended.
Country	Setting this parameter to the appropriate value enables the MAX TNT to generate country-specific local call-progress tones (such as dial tone, busy signals, and so forth), based on the ITU-T specification TSB Circular 18: <i>Update of Supplement No. 2, ITU-T (former CCITT) Blue Book, Fascicle II.2 - Various tones used in national networks</i> . The following country-specific call progress tones are currently supported by MultiVoice: Argentina, Australia, Belgium, China, Costa Rica, Finland, France, Germany, Hong Kong, Italy, Japan, Korea, Mexico, Netherlands, New Zealand, Singapore, Spain, Sweden, Switzerland, United Kingdom, and the United States (the default).

For example, the following commands configure the system in a recommended way for VoIP call handling:

```
admin> read answer-defaults
ANSWER-DEFAULTS read

admin> set session-info idle-timer = 0
```

```
admin> write
ANSWER-DEFAULTS written

admin> read system
SYSTEM read

admin> set max-dialout-time = 60

admin> set parallel-dialing = 32

admin> set country = us

admin> write
SYSTEM written
```

VoIP call management and performance settings

In an H.323 environment, settings in the default VoIP profile are used for processing all VoIP calls. In an SS7 environment, settings in the default VoIP profile are used only for settings that are not superseded by values in IPDC messages.

Most of the VoIP profile settings and IPDC messages are documented in the *MultiVoice for the MAX TNT Configuration Guide*. That guide includes information about call performance settings, such as voice compression, voice packet size, silence thresholds, jitter buffers, and Type of Service (TOS) management. It also documents the H.323v2 call management settings, which include Gatekeeper communication, call signaling and progress tones, dialing options such as single-stage dialing, and PIN or CLID authentication of VoIP calls.

For details about VoIP profile settings that are new in MAX TNT TAOS 8.0.0, see “New VoIP profile settings in MAX TNT TAOS 8.0.0” on page 265. For information about new or modified IPDC messaging, see “IPDC message support for modifying parameters” on page 269.

Real-time fax (T.38)

MultiVoice real-time fax is an implementation of the ITU-T T.38 standard for fax transmission across IP networks, using the VoIP framework for call establishment, fax initiation, and detection of an incoming fax call.

Note: Real-time fax communications require guaranteed quality of service between the two fax-capable Gateways. The packet loss on the network must be less than 1%.

Real-time fax calls begin when a VoIP call is placed from an originating fax machine to the answering machine. If the MAX TNT is configured to perform out-of-band dual tone multifrequency (DTMF) signaling, a DSP automatically enables inband DTMF signaling at the start of the fax call. When the destination fax machine picks up the call and sends an answer tone, known as a CED tone, the destination Gateway detects this tone and initiates a switchover to real-time fax on both itself and the Gateway at the other end of the call. When the switchover is complete, the fax transmission proceeds normally.

You must create the appropriate coverage areas on the MultiVoice Access Manager to ensure that fax calls are routed between Gateways that are fax capable. For details, see the *MultiVoice Access Manager User's Guide* at <http://www.ascend.com/doclibrary>.

Overview of real-time fax settings

Following are the parameters (shown with default values) for enabling and improving the performance of real-time fax processing. Changes to these parameters take effect with the next VoIP call.

```
[in VOIP/{ 0 0 }:rt-fax-options]
rt-fax-enable = no
ecm-enable = yes
low-latency-mode = yes
command-spoof = yes
local-retransmit-lsf = yes
```

Parameter	Specifies
RT-Fax-Enable	Enable/disable T.38 fax call processing. When the parameter is set to <code>no</code> (the default), fax tones are passed as if they were normal voice samples, and the other parameters in the subprofile are not applicable. When the parameter value is set to <code>yes</code> , this MAX TNT switches over from voice session to fax upon detection of a CED tone or V.21 HDLC flag.
ECM-Enable	Enable/disable error correction mode (ECM) for real-time fax calls. When the parameter is set to <code>yes</code> (the default), fax frames can be retransmitted in the event that a frame is not received correctly. ECM frames are relayed end to end between terminals. Setting the parameter to <code>no</code> disables ECM, so fax frames containing errors are not corrected.
Low-Latency-Mode	Enable/disable low latency mode for real-time fax operations over networks with low packet loss and low latency characteristics. Low latency mode allows operation on networks with less than 2.5 seconds or less of aggregate latency between pages. When the parameter is set to <code>no</code> , a minimum of 10 seconds delay is added to processing fax calls to allow interpretation of T.30 frames and implement spoofing.
Command-Spoof	Enable/disable spoofing of certain fax commands. Command spoofing is a method of improving performance and reducing fax errors on low latency networks.
Local-Retransmit-LSF	Enable/disable local retransmission of a low speed fax frame if no response is detected from the destination fax. This is designed to reduce fax transmission errors on low packet loss networks

In an SS7 environment, values in IPDC messages override corresponding call management settings in the default VoIP profile. For information about IPDC support for real-time fax, see “IPDC message support for T.38 fax and transparent modem” on page 270.

Example real-time fax configuration

For example, the following commands enable T.38 fax call processing and leave all performance parameters enabled:

```
admin> read voip { 0 0 }
VOIP/{ 0 0 } read
admin> set rt-fax-options rt-fax-enable = yes
```

```
admin> write
VOIP/{ 0 0 } written
```

Transparent modem

MultiVoice supports a transparent data mode that enables users to run a modem on a VoIP channel, regardless of the audio codec that is in use.

Overview of transparent modem settings

Following is the parameter for enabling the transparent modem features, shown with the default setting:

```
[in VOIP { 0 0 }]
g711-transparent-data = no
```

Parameter	Specifies
G711-Transparent-Data	Enable/disable transparent modem mode. When the parameter is set to yes , when the MAX TNT detects a modem in a VoIP channel, the unit transparently requests end-to-end G.711 encoding and bandwidth for the call, in a process similar to that used by real-time fax. The echo cancelers are disabled when the MAX TNT enters this mode, thus providing transparent G.711 encoding. The data is encoded transparently as an audio-mode type, either G.711 μ -law (64Kbps) or G.711 A-law (64Kbps). Settings take effect with the next incoming PSTN call. A separate license is not required for this feature.

In an SS7 environment, values in IPDC messages override corresponding call management settings in the default VoIP profile. For information about IPDC support for transparent modem, see “IPDC message support for T.38 fax and transparent modem” on page 270.

Example transparent modem configuration

The following commands enable the transparent modem feature on VoIP channels:

```
admin> read voip { 0 0 }
VOIP/{ 0 0 } read

admin> set g711-transparent-data = yes

admin> write
VOIP/{ 0 0 } written
```

Using transparent modem with real-time fax

If the MAX TNT has been licensed for real-time fax, users can run either a high-speed modem with speeds greater than 2400 bps or a fax terminal in the VoIP channel. This capability provides a fallback for real-time fax transmissions. Both fax terminals and high-speed modems transmit a single tone when they answer a call, but each type of equipment uses a different tone. The MAX TNT detects the type of equipment in use on the basis of its answer tone. When it detects the equipment answering the call, the MAX TNT sends H.245 request-mode messages to request a switchover from the current audio codec to either G.711 with no echo canceler (for transparent modem) or T.38 data mode (for real-time fax).

Transparent data is encoded as an audio-mode type, either G.711 μ -law (64Kbps) or G.711 A-law (64Kbps). Real-time fax (if supported) is encoded as data-mode type T.38 fax.

Note: Transparent data mode introduces an H.245 request-mode message that is not backward compatible with the real-time fax feature provided by previous MultiVoice releases. To interoperate with a Gateway using transparent mode, all Gateways must be upgraded to MAX TNT TAOS 8.0.0.

Example real-time fax and transparent modem configuration

The following commands enable both real-time fax and the transparent modem feature for high-speed modems:

```
admin> read voip { 0 0 }
VOIP/{ 0 0 } read

admin> set rt-fax-options rt-fax-enable = yes

admin> set g711-transparent-data = yes

admin> write
VOIP/{ 0 0 } written
```

Limitation for low-speed modems

Real-time fax cannot be used concurrently with low-speed modems (2400bps or less) because these modems use the same answer tone as fax terminals. If a low-speed modem is used on a VoIP channel that is enabled for real-time fax, the Gateway detects a fax answer tone and requests T.38 encoding. The ingress Gateway (typically the Gateway on which the modem call originated) can accept the T.38 encoding request or reject the request, which causes the egress Gateway to terminate the call.

New VoIP profile settings in MAX TNT TAOS 8.0.0

The following parameters (shown with default values) are new or modified in MAX TNT TAOS 8.0.0:

```
[in VOIP/{ 0 0 }]
voice-ann-dir = /current
allow-g711-fallback = yes
allow-coder-fallback = yes
choose-dsp-via = voip-centric
trunk-quiesce-enable = no
early-ringback-enable = no
trunk-prefix-enable = no
```

Parameter	Specifies
Voice-Ann-Dir	Location of voice announcement files on a PCMCIA flash memory card in the MAX TNT unit. In previous releases, the value was read-only. In MAX TNT TAOS 8.0.0, administrators can create directories on the flash memory file system and specify a location for voice announcement files. See “Storing voice announcements in the FAT-16 flash memory file system” on page 266.

Parameter	Specifies
Allow-G711-Fallback	Enable/disable selection of the G.711 codec if the Gateway is unable to select its preferred codec. This parameter does not apply if Allow-Coder-Fallback is set to no. For details, see “Allowing fallback to alternate codecs” on page 267.
Allow-Coder-Fallback	Enable/disable selection of an alternate codec if the Gateway is unable to select its preferred codec. For details, see “Allowing fallback to alternate codecs” on page 267.
Choose-DSP-Via	<i>Not currently supported.</i>
Trunk-Quiesce-Enable	Enable/disable deactivation of a T1 PRI line when a Gateway is unavailable. For details, see “Deactivating trunks used for VoIP calls” on page 267.
Early-Ringback-Enable	Enable/disable generation of an early ringback tone on networks experiencing long call setup times. If the parameter is set to yes, the near-end Gateway plays a ringback tone to the caller as soon as a call connection is established with the far-end Gateway.
Trunk-Prefix-Enable	Enable/disable identification of the entry (ingress) trunk number to the exit (egress) Gateway or call signaling entity by prepending the ingress trunk number to the DNIS number.

Storing voice announcements in the FAT-16 flash memory file system

By default, MultiVoice callers are notified of call progress by DTMF-based tones. The tones report easily recognized call states such as ringback, busy signal, and so forth, as well as tones specific to MultiVoice, such as PIN prompt, which are not as easily recognized by callers. In previous MultiVoice releases, the MAX TNT introduced support for the playback of custom voice announcements to callers to indicate call progress. For details about how voice announcements work, and for information about managing them in the MAX TNT, see the *MultiVoice for the MAX TNT Configuration Guide* at <http://www.ascend.com/doclibrary>.

With MAX TNT TAOS 8.0.0, you can create directories on the flash memory file system and specify a location for voice announcement files. After creating the directory on a flash card and moving voice announcement files into it, specify the pathname in the Voice-Ann-Dir setting. For example, the following commands create a directory named `messages` and a subdirectory named `announce` on the flash card in slot 1:

```
admin> mkdir 1/messages
admin> mkdir 1/messages/announce
```

The following command loads a voice-announcement file named `busy.au` from a TFTP server at 10.10.10.10 to the `/current` directory on flash card 1 (flash card 1 is the default):

```
admin> load file network 10.10.10.10 busy.au
```

The following command moves the `busy.au` file to the new subdirectory on flash card 1:

```
admin> mv 1/current/busy.au 1/messages/announce/busy.au
```

The following commands inform the MultiVoice subsystem of the location of the voice announcement files:

```
admin> read voip { 0 0 }
VOIP/{ 0 0 } read
```

```
admin> set voice-ann-dir = /messages/announce  
  
admin> write  
VOIP/{ 0 0 } written
```

You can specify a pathname up to 40 characters long. When the system receives a request to play an announcement, it looks in the specified directory on the flash card in slot 1. If the card is not present or the voice announcement file is not found, the system looks for the specified directory on flash card 2.

Allowing fallback to alternate codecs

Voice is transmitted across an IP network as compressed audio frames. The Packet-Audio-Mode parameter in the default VoIP profile specifies the preferred audio codec used by the Gateways to compress and uncompress analog speech and digital audio frames.

In MAX TNT TAOS 8.0.0, you can set the following parameters (shown with default values) to specify how the system behaves when the preferred codec is not supported:

```
[in VOIP/{ 0 0 }]  
allow-g711-fallback = yes  
allow-coder-fallback = yes
```

Normally, an H.323 stack advertises a list of supported audio codecs. If the preferred codec is present on a list received from a far-end Gateway, that codec is always selected. Otherwise, the system selects an alternate codec that matches one from its supported list.

The Allow-Coder-Fallback parameter can be set to `no` to override the default system behavior and force the Gateway to reject the call if it is unable to select its preferred codec. If this parameter is set to `no`, the Allow-G711-Fallback parameter has no effect.

If Allow-Coder-Fallback parameter is set to `yes`, you can set the Allow-G711-Fallback parameter to `no` to prevent the system from selecting the G.711 codec when selecting an alternate codec. In this case, the system terminates the call if G.711 is the only available choice and it is not the preferred code. This setting affects VoIP, fax, and transparent modem calls.

Deactivating trunks used for VoIP calls

The trunk deactivation feature enables MultiVoice Gateways to automatically deactivate trunks used for VoIP calls when a Gateway becomes unavailable. This feature allows Gatekeepers in the MultiVoice network to route calls to other available Gateways, to use network resources more efficiently and improve service quality for users.

Note: In this release, only T1 trunks that use ISDN PRI signaling and have been configured for VoIP can be deactivated system-wide by using this feature.

Trunk deactivation prevents the PSTN switch from routing subsequent calls to the trunks configured for VoIP. Current calls remain active until those calls are terminated by the caller or PSTN. When trunk deactivation is enabled, trunks configured to accept VoIP calls are made unavailable to the PSTN under the following conditions:

- A Gateway cannot register with either a primary or secondary Gatekeeper.
- A Gateway's trunk connection with the PSTN is unavailable, so that Gateway is forced to unregister itself from its Gatekeepers.

Previously, when a Gateway could not register with the primary and secondary Gatekeeper, the caller heard a fast busy signal because the PSTN switch continued to route calls to the trunks on that Gateway. Deactivating the trunk changes the trunk state to inform the PSTN switch aware that those trunks are not available.

Previously, when a VoIP call could not connect because a trunk was not operating, the caller heard a fast busy signal, because the Gatekeeper continued to route calls to that Gateway as long as it remained registered. Deactivating the trunk forces the Gateway to unregister from all known Gatekeepers, which causes the Gatekeepers to reroute new calls to other Gateways. When any one of the Gateway's trunks comes back in service, that Gateway starts registering itself with one of its known Gatekeepers. The Gatekeeper then begins to route calls to this Gateway.

The following commands enable trunk deactivation for T1 PRI lines configured for VoIP:

```
admin> read voip { 0 0 }
VOIP/{ 0 0 } read

admin> set trunk-quiesce-enable = yes

admin> write
VOIP/{ 0 0 } written
```

Enabling early ringback

For certain VoIP network configurations, such as satellite IP networks, wireless networks, or networks using channel-associated signaling (CAS) trunks, call setup times can be quite long. Callers might hang up before the call completes because they hear no call progress tones until RTP carries ringback from the far end PSTN. Early ringback allows the MAX TNT to generate a ringback tone locally, as soon as the call is started on the far-end Gateway.

Note: Early ringback is intended for use only on networks that experience long call setup times. Its use for other network configurations is not recommended, and might result in erroneous ring-to-busy and ring-to-failure announcements.

The following commands enable early ringback:

```
admin> read voip { 0 0 }
VOIP/{ 0 0 } read

admin> set early-ringback-enable = yes

admin> write
VOIP/{ 0 0 } written
```

Trunk prefixing

Trunk prefixing enables the MAX TNT to identify the entry (ingress) trunk number to the exit (egress) gateway or call signaling entity by prepending the ingress trunk number to the DNIS number. Trunk groups must be in use system-wide.

When trunk prefixing is enabled, the system obtains the trunk group number of the ingress T1 trunk from the `trunk-group` setting in the T1 line profile, and prepends it to the detected DNIS number. The Q.931 Called Party Number information element (IE) in an H.225/Q.931 SETUP message then contains the DNIS number prefixed by the incoming trunk number. The destination address value of the SETUP user-to-user information element (UUIE) is not currently encoded.

For example, the following commands enable trunk prefixing, beginning with the next VoIP call the MAX TNT receives:

```
admin> read voip { 0 0 }
VOIP/{ 0 0 } read

admin> set trunk-prefix-enable = yes

admin> write
VOIP/{ 0 0 } written
```

IPDC message support for modifying parameters

With MAX TNT TAOS 8.0.0, MAX TNT units provide limited support for IPDC messages used to modify the following values for VoIP calls. The request modify packet pass-through call (RMCP) message (0x0015) and accept modify packet pass-through call (AMCP) message (0x0016) allow modification of the following values for VoIP calls.

- VoIP encoding type (G.711 μ -law, G.711A-law G.729, or G.723).
- Packet loading rate in frames per packet (value depends on VoIP encoding type)
- Source port type (currently, only the SCN value is supported).
- Destination port type (currently, only the RTP value is supported).
- Listen IP address.
- Listen RTP port number.
- Send IP address.
- Send RTP port number.

The MAX TNT can allocate its own system IP address as the listen IP address and RTP port and can specify its own send address and RTP port. For VoIP calls, you must avoid routing RTP traffic through the MAX TNT shelf controller. For that reason, when allowing the MAX TNT Gateway to allocate its own address, you must set the System-IP-Addr parameter in the IP-Global profile to an interface address other than the shelf-controller Ethernet port. For example, the following commands set the system address to the address of a port on an Ethernet card in slot 12:

```
admin> get ip-interface { { 1 12 1 } 0 } ip-address
[in IP-INTERFACE/{ { shelf-1 slot-12 1 } 0 }:ip-address]
ip-address = 1.1.1.1/24

admin> read ip-global
IP-GLOBAL read

admin> set system-ip-addr = 1.1.1.1/24

admin> write
IP-GLOBAL written
```

In addition, you must make sure that VoIP calls can always find a route to the next-hop Gateway on the path to the destination VoIP Gateway. The route can be learned dynamically or configured as a static route. Many sites choose to configure default routes for VoIP traffic, so that RTP packets are never dropped due to lack of routing information. For example, the following commands configure a default route named VoIP to a next-hop Gateway at 2.2.2.2:

```
admin> new ip-route voip
IP-ROUTE/voip read

admin> set gateway = 2.2.2.2/24
```

```
admin> write
IP-ROUTE/VoIP written
```

IPDC message support for T.38 fax and transparent modem

Previously, transparent data for fax and modem calls was available only in an H.323 environment or for IPDC calls running G.711 codecs for VoIP. In this release, IPDC message request packet pass-through call (RCCP), accept packet pass-through call (ACCP), request modify for packet pass-through call (RMCP), and accept modify packet pass-through call (AMCP) messages enable an SS7 signaling gateway to direct the MAX TNT to enter T.38 fax mode or transparent modem mode on the basis of tone detection. In addition, the signaling gateway can control echo cancelation by disabling it or setting it to 32 milliseconds on a per-call basis.

The notify tone (NTN) message is used to notify the signaling gateway when an asynchronous fax or modem tone is detected. The MAX TNT sends this message to the signaling gateway if either fax or modem tone detection is enabled and the unit sees the tone. The MAX TNT detects fax tone if `rt-fax-enable` is set to `yes` in the default VoIP profile or if it receives the relevant IPDC message from the signaling gateway.

The MAX TNT detects modem tone if `g711-transparent-data` is set to `yes` in the default VoIP profile or if it receives the relevant IPDC message from the signaling gateway.

For an introduction to the real-time fax feature, see “Real-time fax (T.38)” on page 262. For an introduction to the transparent modem feature, see “Transparent modem” on page 264.

New trunk features for VoIP calls

With MAX TNT TAOS 8.0.0, MAX TNT units provide a configurable timer for T1 lines that use inband signaling, a true connect feature to avoid charges for VoIP calls, and a calling line ID (CLID) generated by the MultiVoice Access Manager (MVAM).

Configurable interdigit timer for T1 inband signaling

When a T1 line uses inband signaling, you can enable Collect-Incoming-Digits to cause the DSP to decode the calling and called DTMF digits on the line, making DNIS and CLID information available for authentication and accounting. Following is the relevant parameter, shown with a sample setting:

```
[in T1/{ any-shelf any-slot 0 }:line-interface]
collect-incoming-digits = yes
```

In previous releases, when this feature was enabled, the T1 DSP always waited for 3 seconds after collecting the last digit before considering DNIS or automatic number identification (ANI) collection complete. This 3-second timeout slowed down call setup times, and was unnecessary when a switch or PBX was generating the DTMF DNIS/ANI information with digit and interdigit times much smaller than 3 seconds. To improve call setup times, especially for VoIP calls with single-stage-dial, you can now configure the timeout for collecting incoming digits. Following is the relevant parameter, shown with its default value:

```
[in T1/{ any-shelf any-slot 0 }:line-interface]
t1-inter-digit-timeout = 3000
```

Parameter	Specifies
T1-Inter-Digit-Timeout	<p>Number of milliseconds the T1 DSP waits between digits before considering DNIS/ANI collection complete. For backward compatibility, the default is 3 seconds. The valid range is 100 to 6000 milliseconds. The setting takes effect with the next incoming call.</p> <p>Specifying a lower value improves call setup times. This is especially important for VoIP calls with single-stage-dial.</p> <p>This parameter does not apply unless Collect-Incoming-Digits is set to yes.</p>

For example, the following commands specify a timeout of half a second:

```
admin> read t1 { 1 2 3 }
T1/{ shelf-1 slot-2 3 } read
admin> set line-interface collect-incoming-digits = yes
admin> set line-interface t1-inter-digit-timeout = 500
admin> write
T1/{ shelf-1 slot-2 3 } written
```

Delaying charges until call is answered (true connect)

In earlier releases, incoming VoIP calls from the PSTN were connected at the near end Gateway before any H.323 signaling was sent to the far end Gateway. As a result, a PSTN charge was incurred at the time of connection to the near-end Gateway, before the called party received and answered the call from the far-end Gateway.

Now, you can change this behavior by enabling true connect. When this feature is enabled, alerting and connect messages sent to the PSTN switch are delayed to match the equivalent H.323 signaling to avoid incurring charges before a VoIP call has been answered.

The true connect feature requires a default call type of VoIP on T1 or E1 trunks accepting incoming VoIP calls. Following are the relevant parameters, shown with sample settings:

```
[in VOIP { 0 0 }]
true-connect-enable = yes

[in T1/{ shelf-1 slot-10 1 }:line-interface]
default-call-type = voip

[in E1/{ shelf-1 slot-11 1 }:line-interface]
default-call-type = voip
```

Parameter	Specifies
True-Connect-Enable	Enable/disable delay of PSTN alerting and connect messages to match the equivalent H.323 alerting and connect messages. The default setting is <code>no</code> , which results in the caller incurring a PSTN charge at the time of connection to the near-end Gateway, before the called party has received and answered the call from the far end Gateway. If set to <code>yes</code> , an alerting message is sent to the ingress PSTN switch only when an H.323 alerting message is received on the ingress VoIP Gateway. Similarly, a PSTN connect message is sent only when the H.323 VoIP call has been answered. This ensures that no charges are incurred for incomplete calls. The setting takes effect with the next incoming call. It has no effect on outbound calls.
Default-Call-Type	Must be set to VoIP for T1 or E1 trunks with incoming VoIP calls that require true connect. Note that setting this parameter to VoIP causes <i>all</i> calls received on the trunk to be mapped to VoIP.

For example, the following commands enable delayed PSTN alerting and connect messages on trunk lines configured with a default VoIP call type:

```
admin> read voip { 0 0 }
VoIP { 0 0 } read

admin> set true-connect-enable = yes

admin> write
VoIP { 0 0 } written
```

Note: For ISDN trunks, Lucent recommends that you set the T310 timer on the telephone company switch or PBX to 30 seconds or greater when using the true connect feature. because the T310 timeout value includes the time that the called party's telephone is ringing, a 10-second timeout can cause the near-end Gateway to disconnect the call too soon.

When the true connect feature is enabled and a VoIP call fails before the PSTN call is fully connected, the Gateway is still able to send an appropriate tone or voice announcement to the caller.

Gatekeeper CLID substitution

When MultiVoice Gateways are connecting VoIP calls, they can transmit a calling line ID (CLID) generated by the MVAM software on the Gatekeeper instead of the PSTN-generated CLID collected on the trunk line. CLID substitution allows the MultiVoice network to provide the appropriate E.164 address for both the called and calling telephone numbers to the respective PSTN, and for use by external applications.

In certain configurations in which the Gateways connecting the call reside in different area codes or countries, the CLID received from the PSTN must be changed to provide the appropriate calling number information to the local carrier, or to call management and billing applications.

When the MVAM receives the CLID from a Gateway, it translates the CLID to the appropriate dial string, adding or removing country codes and area codes as appropriate for the respective locations of the callers. The Gatekeeper then reports the revised CLID to the Gateways as part of the admission confirmed (ACF) message.

RT-24 (proprietary) codec support

The RT-24 codec is a Lucent Technologies proprietary audio codec that compresses speech samples from 64Kbps pulse code modulation (PCM) to 2.4Kbps, reducing the effective bandwidth required for transmission across the IP network.

This codec uses a 22.5-millisecond audio frame, and encapsulates audio at 8 bytes per frame. The decoder produces 180 samples of audio from the 8-byte encoder output. The RT-24 codec is available for both H.323 VoIP calls and SS7 VoIP calls.

When the RT-24 codec is selected, the MultiVoice Gateway attempts to determine if that codec is supported by the other Gateway during H.245 capability negotiation. If both sides agree to use RT-24 as the preferred codec, both Gateways enable RT-24 on the allocated DSPs to compress and decompress audio after the H.245 open logical channel message is exchanged.

Note: RT-24 is a Lucent Technologies proprietary codec, which is available only on MultiVoice Gateways running MAX TNT TAOS 8.0.0. MultiVoice cannot use this codec when communicating with a third-party VoIP gateway.

To enable RT-24 audio processing, set the packet-audio-mode parameter in the default VoIP profile to the selected codec as illustrated by the following:

```
admin> read voip { 0 0 }
VOIP/{ 0 0 } read
admin> set packet-audio-mode = rt24
admin> write
VOIP/{ 0 0 } written
```

G.728 codec support

G.728 is a Low-Delay Code Excited Linear Prediction (LD-CELP) based audio codec that provides toll-quality audio at a bit-rate of 16Kbps. With a frame size of only 2.5 milliseconds, G.728 also has a very low delay. Although the MultiVoice implementation of G.728 uses a frame size of 5 milliseconds, the bitstream from the audio codec is the same as described in the ITU-T standard and can thus be decoded by any G.728 decoder.

Each MultiDSP card supports a maximum of 48 simultaneous G.728 calls for both H.323 VoIP and SS7 VoIP call processing.

When the G.728 codec is selected, the MultiVoice Gateway attempts to determine if the G.728 codec is supported by the other Gateway during H.245 capability negotiation. If both sides agree to use G.728 as the preferred codec, both Gateways use G.728 to compress and decompress audio after the H.245 open logical channel message is exchanged.

Note: Although MultiVoice uses a 5-millisecond frame for G.728 processing, it is compatible with any third-party G.728 decoder. However, if a MultiVoice Gateway attempts to communicate with a third-party VoIP gateway transmitting an odd number of 2.5 millisecond frames per IP packet, the call will fail.

When you enable G.728 audio processing in this release the Silence-Det-Cng parameter must be set to no (its default value). The following commands enable G.728 processing:

```
admin> read voip { 0 0 }
VOIP/{ 0 0 } read
admin> set packet-audio-mode = g728
```

```

admin> set silence-det-cng = no
admin> write
VOIP/{ 0 0 } written

```

SNMP: Support for the VoIP MIB (ascend 28)

The VoIP MIB enables network management stations to monitor MultiVoice Gateway operations using SNMP. Attributes in the MIB can be obtained by SNMP Get and Get-Next operations. The MIB uses the following object identifiers for identifying MultiVoice Gateway or Gatekeepers to a network manager:

- voipCfgGroup (voipGroup 1)
- voipCfgGkGroup (voipCfgGroup 1)
- voipCfgGwGroup (voipCfgGroup 2)

The MIB uses the following tables for identifying MultiVoice Gatekeeper and Gateway functions.

```

voipCfgGkTable OBJECT-TYPE (voipCfgGkGroup 1)
    SYNTAX SEQUENCE OF VoipCfgGkEntry
    ACCESS not-accessible
    STATUS mandatory
    DESCRIPTION A list of entries for H323 network Gatekeeper.

voipCfgGkEntry OBJECT-TYPE (voipCfgGkTable 1)
    SYNTAX VoipCfgGkEntry
    ACCESS not-accessible
    STATUS mandatory
    DESCRIPTION An entry holding information about the Gatekeeper for
    the system.
    INDEX (voipCfgGkIndex)

VoipCfgGkEntry:
    SEQUENCE :
        voipCfgGkIndex-INTEGER
        voipCfgGkStatus-INTEGER
        voipCfgGkIpAddress-IpAddress)

voipCfgGkIndex OBJECT-TYPE ( voipCfgGkEntry 1)
    SYNTAX INTEGER
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION This number uniquely identifies the Gatekeeper.

voipCfgGkStatus OBJECT-TYPE (voipCfgGkEntry 2)
    SYNTAX INTEGER:
        registered(1)
        not_registered(2)
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION This value indicates whether the gateway is registered
    with the Gatekeeper.

voipCfgGkIpAddress OBJECT-TYPE (voipCfgGkEntry 3)
    SYNTAX IpAddress
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION The IP address of the Gatekeeper.

```

```
voipCfgGwVpnMode OBJECT-TYPE (voipCfgGwGroup 1)
    SYNTAX INTEGER:
        no (1)
        yes(2)
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION Virtual Private Network Toggle Switch.

voipCfgGwPktAudioMode OBJECT-TYPE (voipCfgGwGroup 2)
    SYNTAX INTEGER:
        other(1)
        g711_ulaw(2)
        g711_alaw(3)
        g723(4)
        g729(5)
        g723_6_4kps(6)
    ACCESS read-only
    STATUS mandatory
    DESCRIPTION Audio Coder to be used for voice packetization.
```

The `voipCfgGwVpnMode` and `voipCfgGwPktAudioMode` objects can be accessed using index 0 because they are separate leaves in the MIB tree.

The `voipCfgGkIndex`, `voipCfgGkCurrent` and `voipCfgGkIpAddress` objects are located in the `voipCfgGkTable` table. They can be obtained using `voipCfgGkIndex` as an index.

SNMP: Traps for VoIP-related conditions

With MAX TNT TAOS 8.0.0, VoIP-enabled MAX TNT units can generate traps for the following MultiVoice Gateway events:

- Change in the call logging server
- Change in configured Gatekeeper for VoIP
- Change in state of a WAN line

For the traps to be sent, traps must be enabled in the system and the individual trap conditions must be set to `yes`. For details about enabling traps, see the *MAX TNT Administration Guide*. Following are the relevant parameters (shown with default values) for enabling the individual trap conditions:

```
[in TRAP/""]
call-log-serv-change-enabled = no
voip-gk-change-enabled = no
wan-line-state-change-enabled = no
```

Parameter	Specifies
Call-Log-Serv-Change-Enabled	<p>Enable/disable trap generation when the call-logging server changes. If the call-logging server index is changed or if the IP address of the active call-logging server is changed, this trap sends the following information to the SNMP manager:</p> <ul style="list-style-type: none"> • The new call logging server index (callLoggingServerIndex) • The IP address of new call logging server (callLoggingServerIPAddress) • The absolute time to show when the server change occurred (sysAbsoluteCurrentTime) (Ascend Trap 38)
Voip-GK-Change-Enabled	<p>Enable/disable trap generation when the registered Gatekeeper changes. If a new Gatekeeper is registered with the Gateway, a register request (RRQ) message is sent from the Gateway to the new Gatekeeper. When the Gateway receives the admission request (ARQ) message from the new Gatekeeper, this trap sends the following information to the SNMP manager:</p> <ul style="list-style-type: none"> • The new Gatekeeper index (voipCfgGkIndex) • The IP address of new Gatekeeper (voipCfgGkIpAddress) • The absolute time to show when the Gatekeeper change occurred (sysAbsoluteCurrentTime) (Ascend Trap 39)
WAN-Line-State-Change-Enabled	<p>Enable/disable trap generation if the state of an E1 or T1 line changes. This trap sends the following information to the SNMP manager:</p> <ul style="list-style-type: none"> • The T1 or E1 line interface index (wanLineIfIndex) • The line usage (wanLineUsage). This usage is reported as trunk, quiesced, or disabled. • The absolute time to show when the line state changed (sysAbsoluteCurrentTime) (Ascend Trap 40)

NavisAccess support for VoIP call reporting

MAX TNT TAOS 8.0.0 supports basic VoIP call reporting using NavisAccess. This includes the capability to generate Start records, Stop records, and Call Progress records for both VoIP and fax calls. These records allow NavisAccess to monitor Gateway resource usage and provide information to create billing records. Each VoIP call can generate two or more records.

Start records

A Start record reports the point in the call where a speech communications is established. Start records can provide the following information:

Attribute	Specifies
Ascend-Call-Direction	Direction of the call between the Gateway and PSTN. The reported values are Ascend-Call-Direction-Incoming (0) and Ascend-Call-Direction-Outgoing (1). (Ascend Trap 48)
NAS-Port	Encoded NAS port used for this call. (RFC Trap 5)

Attribute	Specifies
NAS-Port-Type	Encoded NAS port used for this call. The value 7 for this attribute identifies a VoIP call. (RFC Trap 61)
NAS-IP-Address	NAS IP address associated with this call. (RFC Trap 4)
Session-Id	NAS session index recorded in the session table for this call. (RFC Trap 44)
Ascend-Modem-PortNo	DSP/modem port allocated for processing this call. This value is part of the resource count information, and is repeated each time it is allocated for a call. (Ascend Trap 120)
Ascend-Modem-SlotNo	Slot where the DSP/modem card associated with the reported Ascend-Modem-PortNo is located. This value is part of the resource count information, and is repeated each time it is allocated for a call. (Ascend Trap 121)
Ascend-Modem-ShelfNo	Shelf where DSP/modem card allocated for processing this call is installed. This is part of the resource count information, and is repeated each time it is allocated for a call. (Ascend Trap 122)
Called-Station-Id (DNIS)	Dialed number string reported by the Gateway for the called destination. (RFC Trap 30)
Ascend-Dialed-Number	Dialed number string used by the Gateway to complete the call. (Ascend Trap 24)
Service-Type	Requested type of service, the value of the Type of Service byte, for this call. (RFC Trap 6)
Ascend-H323-Destination-NAS-ID	NAS IP address used to route the call to the connecting Gateway. (Ascend Trap 22)
Ascend-H323-Gatekeeper-IP	IP address of the Gatekeeper used to route the call. The Gateway is registered with this Gatekeeper. (Ascend Trap 19)
Ascend-Global-Call-Id	IP address used by the Gatekeeper to identify the connecting Gateway for this call. (Ascend Trap 20)
Ascend-H323-Conference-ID	IP address used to identify the called destination. (Ascend Trap 21)
Ascend-H323-Pre-session-Time	Time from the moment the caller finishes dialing the destination telephone number until the moment the speech path is established to the called destination. (Ascend Trap 198)
Ascend-H323-Dialed-Time	Time the user spends dialing the destination telephone number. This value will be zero for call originating from the LAN. (Ascend Trap 23)
Ascend-Session-Type	Audio codec used for processing the call. (Ascend Trap 18)

Stop records

A Stop record is generated at the moment when MultiVoice begins to tear down the speech path or when an incoming call to a Gateway fails to connect. A Start record can contain following information:

Attribute	Specifies
Acct-Session-Time	Time from the moment the speech path is established to the called destination until the moment MultiVoice begins to tear down the speech path. (RFC Trap 46)
Ascend-Connect-Progress	A number that represents the call connect state at the time the call was terminated. (Ascend Trap 195)
Ascend-Disconnect-Cause	A number that reports the H.323 call disconnection reason. (Ascend Trap 196)
Ascend-H323-Inter-Arrival-Jitter	Estimated interarrival jitter for voice packets received by a Gateway. (Ascend Trap 25)
Ascend-Dropped-Octets	The number of voice frames (in bytes) dropped by a Gateway during call processing. (Ascend Trap 26)
Ascend-Dropped-Packets	Number of voice packets dropped by a Gateway during call processing. (Ascend Trap 26)
Acct-Input-Octets	Number of voice frames (in bytes) received by a Gateway during this call. (RFC Trap 42)
Acct-Input-Packets	Number of voice packets received by a Gateway during this call. (RFC Trap 47)
Acct-Output-Octets	Number of voice frames (in bytes) sent by a Gateway during this call. (RFC Trap 43)
Acct-Output-Packets	Number of voice packets sent by a Gateway during this call. (RFC Trap 48)

Call Progress records

A Call Progress record can be generated during a VoIP call when a change in resource occurs for a fax or transparent modem call. For fax calls, this record includes the modem speed and modulation. A progress message contains all the information included in a Start record.

Tunneling

IPSec for L2TP tunnels and TCP-Clear

With MAX TNT TAOS 8.0.0, the MAX TNT supports the IPSec Authentication Header (AH) and Encapsulating Security Payload (ESP) protocols for secure transmission of IP data packets. Each protocol supports two modes of use: transport mode and tunnel mode. For complete information about the IPSec protocols, see RFC 2401, *Security Architecture for the Internet Protocol* (November 1998), RFC 2402, *IP Authentication Header* (November 1998), and RFC 2406, *IP Encapsulating Security Payload (ESP)* (November 1998).

IPSec security protocols

IPSec AH uses a shared secret (a *key*) to run portions of a data packet through digest algorithms to create a digital fingerprint. The receiving system performs the same process and compares the fingerprints. If the fingerprints match, the receiving system is assured that the packet was sent by the right source and was not altered in transit.

IPSec ESP performs full encryption of the data portion of every packet. The receiving system decrypts the packets before routing them. The encryption/decryption provides the added assurance that packet contents have not been viewed while the packet was in transit.

IPSec encapsulation modes

The MAX TNT supports IPSec transport mode and tunnel mode.

Transport mode operates between two hosts. Transport mode provides security services for higher-layer protocols, which can include selected portions of the IP header and other selected options.

Tunnel mode is required for connections between a host that does not perform IPSec processing and a security gateway. In tunnel mode, IP packets are encapsulated in an outer IP header that specifies the IPSec processing destination (IP-in-IP encapsulation).

Applying IPSec to a tunnel server or TCP connection

IPSec profiles specify an IPSec endpoint as well as the IPSec transforms to use on the data stream transmitted to and from that endpoint. Both endpoints must have matching configurations. (The settings in the *send* configuration of one system must match those in the *receive* configuration of the other system, and vice versa.)

In an IPSec profile, the following parameters (shown with default values) enable the profile, and specify the encapsulation mode and far-end IPSec endpoint address:

```
[in IPSEC/""]
name* = ""
active = no
encap-mode = transport
tunnel-address = 0.0.0.0
```

Parameter	Specifies
Name	Name of the IPSec profile (up to 23 characters).
Active	Enable/disable the profile for use.
Encap-Mode	Encapsulation mode in which IPSec operates. For background information, see “IPSec encapsulation modes” on page 279. The default value is <code>Transport</code> (transport mode). If the parameter is set to <code>tunnel</code> , the data stream is tunneled using IP-in-IP encapsulation. Tunnel mode is required if the IPSec endpoint addresses differ from the TCP endpoint addresses for TCP-Clear sessions. If the parameter is set to <code>optimized</code> , the system uses transport mode if possible (transport mode is more efficient) and uses tunnel mode only when it is required for a particular connection.
Tunnel-Address	IP address of the far-end IPSec endpoint. For an L2TP connection, this is the IP address of the L2TP network server (LNS) at the far end of the tunnel. For a TCP-Clear connection, it is the address of a security gateway or dial-in host.

For example, the following commands configure specify that IPSec processing on the data stream transmitted to and from 1.1.1.1 operates in transport mode:

```
admin> new ipsec securegw
IPSEC/l2tpl-ipsec read

admin> set active = yes

admin> set encap-mode = transport

admin> set tunnel-address = 1.1.1.1

admin> write
IPSEC/l2tpl-ipsec written
```

Note: The example commands above do not create a usable IPSec profile. IPSec AH or IPSec ESP (or both) must be configured for the IPSec profile to have an effect. For details, see “Configuring an IPSec profile for IPSec AH” on page 282 and “Configuring an IPSec profile for IPSec ESP” on page 284.

Applying an IPSec profile to an LNS


The following parameters (shown with default values) associate an IPSec profile with a particular LNS:

```
[in TUNNEL-SERVER/" "]
server-endpoint = ""
ipsec-profile = ""
```

Parameter	Specifies
Server-Endpoint	Name or IP address of the tunnel endpoint. This parameter must specify the same host as the Tunnel-Address parameter in the IPSec profile to be applied.
IPSec-Profile	Name of the IPSec profile (up to 23 characters) that defines the transforms and endpoints for IPSec operations on traffic crossing L2TP tunnels to the specified endpoint. If the Tunnel-Server profile does not specify an IPSec profile name, a normal, nonsecure UDP socket is assigned to the L2TP session. If an IPSec profile name is specified, a new UDP socket is opened and assigned the specified profile settings.

Figure 43 shows a MAX TNT unit operating as an L2TP access concentrator (LAC) with a MAX unit operating as the LNS:

Figure 43. Secure IPSec L2TP tunneling configuration



The following commands apply an IPSec profile named `securegw` to the specified LNS:

```
admin> new tunnel-server
TUNNEL-SERVER/" " read

admin> set server-endpoint = 1.1.1.1

admin> set ipsec-profile = securegw
```



```
admin> write
TUNNEL-SERVER/1.1.1.1 written
```

For details about setting up L2TP configurations, see the *MAX TNT Network Configuration Guide*.

Applying an IPsec profile to a TCP-Clear session

RADIUS profiles use the following attribute-value pair to apply an IPsec profile to a TCP-Clear session:

RADIUS attribute	Value
Ascend-IPSEC-Profile (73)	Name of an IPsec profile that describes the IPsec transforms and endpoints to use for this connection (a string value).

To use these attributes, the RADIUS server must support vendor-specific attributes (VSAs) and the MAX TNT must be configured in VSA compatibility mode. Following are the relevant settings:

```
[in EXTERNAL-AUTH]
auth-type = radius


[in EXTERNAL-AUTH:rad-auth-client]
auth-radius-compat = vendor-specific
```

For details about these settings, see the *MAX TNT Reference Guide*. For details about setting up TCP-Clear connections, see the *MAX TNT Network Configuration Guide*.

In Figure 44, the IPsec endpoints are a MAX TNT unit and Pipeline 220 unit. The TCP endpoint is a TCP host dialing into the MAX TNT. Because the IPsec endpoints are different from the TCP endpoint, the IPsec profile for this connection must specify tunnel mode, as described in “IPsec encapsulation modes” on page 279. For example:

```
admin> get ipsec securegw encap-mode
[in IPSEC/securegw:encap-mode]
encap-mode = tunnel
```

Figure 44. IPsec tunnel mode for TCP-Clear between gateways



The following sample RADIUS profile enables the dial-in host to establish a secure TCP-Clear session to a login host at 1.1.1.5:


```
tcpapp-user Password = "localpw"
  Service-Type = Login,
  Login-Service = TCP-Clear,
  Login-Host = 1.1.1.5,
  Login-TCP-Port = 23,
  Ascend-IPSEC-Profile = securegw
```

The MAX TNT applies an IPsec profile named `securegw` to the session's data stream. The Pipeline 220 must have a corresponding IPsec configuration, and the two IPsec endpoints require Connection or RADIUS profiles for the link between them.

In Figure 45, the dial-in host is running a TCP application that is capable of performing IPsec encapsulation and decapsulation. In this case, the IPsec endpoints and the TCP-Clear endpoint are the same, so the IPsec profile for this connection specifies transport mode. For example:

```
admin> get ipsec dialin encap-mode
[ in IPSEC/dialin:encap-mode ]
encap-mode = transport
```

Figure 45. IPsec transport mode for TCP-Clear with dial-in host



The following sample profile enables the dial-in user to establish a secure TCP-Clear session to a login host at 2.2.2.3 using an IPsec profile named `dialin`.

```
dialin-user Password = "my-password"
  Service-Type = Login,
  Login-Service = TCP-Clear,
  Login-Host = 2.2.2.3,
  Login-TCP-Port = 23,
  Ascend-IPSEC-Profile = dialin
```

Configuring an IPsec profile for IPsec AH

For IPsec AH to operate, you at both ends of the L2TP tunnel must specify a security parameter index (SPI) number, the type of transform to use, and a shared secret (a key). These settings must match in the LAC and LNS configurations. In addition, you can choose to enable *replay protection*, which is used to counter denial-of-service attacks.

Overview of the IPsec AH settings

Administrators at both ends of the tunnel must specify matching IPsec AH configurations. Following are the relevant MAX TNT parameters, shown with their default settings:

```
[ in IPSEC/"" :send-ah ]
active = no
spi = 1
ah-type = none
key =
replay-protection = no

[ in IPSEC/"" :recv-ah ]
active = no
spi = 1
ah-type = none
key =
replay-protection = no
```

Parameter	Specifies
Active	Enable/disable IPSec AH processing for packets sent or received through the tunnel.
SPI	Security parameters index: a 32-bit numeric value from 1 to 2147483647. The SPI in the Send-AH subprofile must match the LNS SPI in its receiving AH configuration, and vice versa. If the LNS is a TAOS unit (such as a MAX unit), the administrator of that unit can use the Secure Connect Manager (SCM) to create and download IPSec configurations in Firewall profiles. The SPI values in SCM are in hexadecimal, while the MAX TNT SPI values are in decimal. You can enter the SPI value here in hexadecimal by preceding the value with 0x. However, the number is still displayed in decimal in the MAX TNT interface.
AH-Type	Type of authentication transform to use. Following are valid values: None (the default): No authentication MD5: MD5 mode, described in RFC 1828 SHA1: SHA1 mode, described in RFC 1852 on the secure hash algorithm (SHA) MD5-HMAC: Version-2 MD5, currently in draft SHA1-HMAC: Version-2 SHA1, currently in draft
Key	An authentication key for hashing: A 64-byte text string that exactly matches the key specified in the LNS IPSec AH configuration.
Replay-Protection	Enable/disable sequence number processing. The receiving system uses a sequence number to detect arrival of duplicate packets within a constrained window. If enabled in the Send-AH subprofile, the MAX TNT generates a sequence number for packets it sends through the tunnel. In this release, the MAX TNT does not verify the sequence of packets it receives from the LNS, even if Replay-Protection is enabled in the Recv-AH subprofile.

Example of an IPSec AH configuration

In the following example, an administrator creates an IPSec profile applying IPSec AH to all data sent and received through an L2TP tunnel to an LNS at the IP address 1.1.1.1:

```
admin> new ipsec l2tp1-ipsec
IPSEC/l2tp1-ipsec read
admin> set active = yes
admin> set encap-mode = transport
admin> set tunnel-address = 1.1.1.1
```

In the next commands, the MAX TNT send configuration must match corresponding parameters in the LNS system's IPSec receive configuration, and vice versa:

```
admin> set send-ah active = yes
admin> set send-ah spi = 43981
admin> set send-ah ah-type = md5
admin> set send-ah key = 4142434445464748494A4B4C4D4E4F50
```

```
admin> set recv-ah active = yes
admin> set recv-ah spi = 43981
admin> set recv-ah ah-type = md5
admin> set recv-ah key = 4142434445464748494A4B4C4D4E4F50
admin> write
IPSEC/l2tpl-ipsec written
```

The next commands apply the IPSec profile to the LNS:

```
admin> read tunnel-server 1.1.1.1
TUNNEL-SERVER/1.1.1.1 read
admin> set ipsec-profile = l2tpl-ipsec
admin> write
TUNNEL-SERVER/1.1.1.1 written
```

Configuring an IPSec profile for IPSec ESP

IPSec ESP provides data encryption as well as replay protection and authentication. If you are configuring ESP in addition to IPSec AH, you can specify encryption alone and rely on AH to provide authentication and replay protection service. If you are specifying IPSec ESP without a corresponding AH configuration, you must include integrity and authentication settings to prevent attacks that could otherwise compromise the security provided by encryption alone.

Overview of IPSec ESP settings

Administrators at both IPSec endpoints must specify matching IPSec ESP configurations. Following are the relevant parameters, shown with their default settings:

```
[in IPSEC/":send-esp]
active = no
spi = 1
version = 0
esp-type = none
iv-len = 32
key =
key2 =
key3 =
auth-type = none
auth-key =
replay-protection = no

[in IPSEC/":recv-esp]
active = no
spi = 1
version = 0
esp-type = none
iv-len = 32
key =
key2 =
key3 =
auth-type = none
auth-key =
replay-protection = no
```

Parameter	Specifies
Active	Enable/disable IPSec ESP processing for packets sent or received through the tunnel.
SPI	Security parameters index: a 32-bit numeric value from 1 to 2147483647. The SPI in the Send-ESP subprofile must match the LNS SPI in its receiving ESP configuration, and vice versa. If the LNS is a TAOS unit (such as a MAX unit), the administrator of that unit can use the Secure Connect Manager (SCM) to create and download IPSec configurations in Firewall profiles. The SPI values in SCM are in hexadecimal, while the MAX TNT SPI values are in decimal. You can enter the SPI value here in hexadecimal by preceding the value with 0x. However, the number is still displayed in decimal in the MAX TNT interface.
Version	ESP version (version 1 or version 2).
ESP-Type	Type of ESP transform to use to encrypt the data portion of IP packets. Following are valid values: None (the default): No encryption DES-CBC: DES-CBC mode, described in RFC 1829, on the US Data Encryption Standard cipher block chaining algorithm 3DES-CBC: 3DES-CBC mode, described in RFC 1851 on the Triple DES-CBC algorithm 40DES-CBC: DES-CBC mode restricted to 40 bits
IV-Len	Number of bits in the Initialization Vector. For ESP-v1, you can set it to 32 (a 32-bit vector) or 64 (a 64-bit vector). For ESP-v2, IV-Len is set to 64 automatically.
Key	An authentication key for ESP: A 16-byte text string that exactly matches the key specified in the LNS IPSec ESP configuration.
Key2	A second 16-byte authentication key, to be used for the second pass of 3DES-CBC mode encryption.
Key3	A third 16-byte authentication key, to be used for the third pass of 3DES-CBC mode encryption.
Auth-Type	Type of authentication transform to use when ESP-v2 is in use. Following are valid values: None (the default): No authentication MD5: MD5 mode, described in RFC 1828 SHA1: SHA1 mode, described in RFC 1852 MD5-HMAC: Version-2 MD5, currently in draft SHA1-HMAC: Version-2 SHA1, currently in draft
Auth-Key	An authentication key to use when ESP-v2 is in use: A 64-byte text string exactly matches the key specified in the LNS IPSec ESP-v2 configuration. This setting does not apply if Version is set to 1.

Parameter	Specifies
Replay-Protection	Enable/disable sequence number processing. The receiving system uses a sequence number to detect the arrival of duplicate packets within a constrained window. If enabled in the Send-AH subprofile, the MAX TNT generates a sequence number for packets it sends through the tunnel. In this release, the MAX TNT does not verify the sequence of packets it receives from the LNS, even if Replay-Protection is enabled in the Recv-AH subprofile.

Example of an IPsec ESP configuration for L2TP

In the following example, an administrator creates an IPsec profile applying IPsec ESP and partial sequence integrity (replay protection) to packets tunneled to and from an LNS at the IP address 1.1.1.1:

```
admin> new ipsec l2tp1-ipsec
IPSEC/l2tp1-ipsec read
admin> set active = yes
admin> set encap-mode = transport
admin> set tunnel-address = 1.1.1.1
```

In the next commands, the MAX TNT send configuration must match corresponding parameters in the LNS system's IPsec receive configuration, and vice versa:

```
admin> set send-esp active = yes
admin> set send-esp spi = 26990
admin> set send-esp version = 2
admin> set send-esp esp-type = des-cbc
admin> set send-esp key = 61083D2A76D57ABC
admin> set send-esp esp-version = 2
admin> set send-esp replay-protection = yes
admin> set recv-esp active = yes
admin> set recv-esp spi = 26990
admin> set recv-esp version = 2
admin> set recv-esp esp-type = des-cbc
admin> set recv-esp key = 61083D2A76D57ABC
admin> set recv-esp esp-version = 2
admin> set recv-esp replay-protection = yes
admin> write
IPSEC/l2tp1-ipsec written
```

The next commands apply the IPsec profile to the LNS:

```
admin> read tunnel-server 1.1.1.1
TUNNEL-SERVER/1.1.1.1 read
admin> set ipsec-profile = l2tp1-ipsec
admin> write
TUNNEL-SERVER/1.1.1.1 written
```

Example of an IPSec ESP configuration for TCP-Clear

In the following example, an administrator creates an IPSec profile applying IPSec ESP to packets tunneled to and from an IPSec security gateway at the IP address 2.2.2.2:

```
admin> new ipsec securegw-1
IPSEC/securegw-1 read
admin> set active = yes
admin> set encap-mode = tunnel
admin> set tunnel-address = 2.2.2.2
```

In the next commands, the MAX TNT send configuration must match corresponding parameters in the far-end security gateway's IPSec receive configuration, and vice versa:

```
admin> set send-esp active = yes
admin> set send-esp spi = 26990
admin> set send-esp version = 2
admin> set send-esp esp-type = des-cbc
admin> set send-esp key = 61083D2A76D57ABC
admin> set send-esp esp-version = 2
admin> set recv-esp active = yes
admin> set recv-esp spi = 26990
admin> set recv-esp version = 2
admin> set recv-esp esp-type = des-cbc
admin> set recv-esp key = 61083D2A76D57ABC
admin> set recv-esp esp-version = 2
admin> write
IPSEC/securegw-1 written
```

Following are sample RADIUS profiles that refer to the IPSec profile:

```
tcpapp1 Password = "secret-1"
    Service-Type = Login,
    Login-Service = TCP-Clear,
    Login-Host = 10.10.10.1,
    Login-TCP-Port = 23,
    Login-Host = 10.10.10.2,
    Login-TCP-Port = 125,
    Ascend-IPSEC-Profile = securegw-1

tcpapp2 Password = "secret-2"
    Service-Type = Login,
    Login-Service = TCP-Clear,
    Login-Host = 10.10.10.1,
    Login-TCP-Port = 23,
    Login-Host = 10.10.10.2,
    Login-TCP-Port = 125,
    Ascend-IPSEC-Profile = securegw-1

tcpapp3 Password = "secret-3"
    Service-Type = Login,
    Login-Service = TCP-Clear,
    Login-Host = 10.10.10.1,
```

```
Login-TCP-Port = 23,  
Login-Host = 10.10.10.2,  
Login-TCP-Port = 125,  
Ascend-IPSEC-Profile = securegw-1
```

L2TP timer options

With MAX TNT TAOS 8.0.0, you can configure new L2TP-related timer parameters. The parameters are located in the L2TP-Config subprofile, which has been added to L2-Tunnel-Global. Use these new parameters when the MAX TNT operates as an L2TP access concentrator (LAC).

Note: The MAX TNT currently operates as a LAC only. It receives incoming PPP calls and initiates a connection to an L2TP network server (LNS). A MAX unit can function as the LNS.

Control-Connect-Establish-Timer

Description: Specifies the maximum number of seconds during which the MAX TNT unit can establish an L2TP tunnel with another unit. Any change you make to this parameter takes effect when the previous timer expires.

Usage: Enter a decimal number from 0 to 600. The default is 60.

Example: `set control-connect-establish-timer = 60`

Dependencies: Control-Connect-Establish-Timer applies only if you have set L2TP-Mode to LAC. The MAX TNT can operate as LAC only.

Location: L2-Tunnel-Global

See Also: First-Retry-Timer, Hello-Timer, L2TP-Mode, LAC-Incoming-Call-Timer, Retry-Count

First-Retry-Timer

Description: Specifies, in milliseconds, the initial interval that the MAX TNT unit waits before making a second attempt to establish an L2TP tunnel with another unit. Any change you make to this parameter takes effect when the previous timer expires.

Usage: Enter a decimal number from 100 to 5000. The default is 1000.

Example: `set first-retry-timer = 1000`

Dependencies: First-Retry-Timer applies only if you have set L2TP-Mode to LAC. The MAX TNT can operate as a LAC only.

Location: L2-Tunnel-Global

See Also: Control-Connect-Establish-Timer, Hello-Timer, L2TP-Mode, LAC-Incoming-Call-Timer, Retry-Count

Hello-Timer

Description: Specifies the interval, in seconds, between Hello messages that the MAX TNT unit sends to another unit. Any change you make to this parameter takes effect when the previous timer expires.

Usage: Specify a decimal number from 0 to 600. The default is 60. 0 specifies that the MAX TNT unit sends no Hello messages.

Example: `set hello-timer = 60`

Dependencies: Hello-Timer applies only if you have set L2TP-Mode to LAC. The MAX TNT can operate as LAC only.

Location: L2-Tunnel-Global

See Also: Control-Connect-Establish-Timer, First-Retry-Timer, L2TP-Mode, LAC-Incoming-Call-Timer, Retry-Count

LAC-Incoming-Call-Timer

Description: Specifies the number of seconds that the MAX TNT unit waits for call setup to complete. Any change you make to this parameter takes effect when the previous timer expires.

Usage: Specify a decimal number from 1 to 600. 60 is the default.

Example: `set lac-incoming-call-timer = 60`

Dependencies: LAC-Incoming-Call-Timer applies only if you have set L2TP-Mode to LAC. The MAX TNT can operate as a LAC only.

Location: L2-Tunnel-Global

See Also: Control-Connect-Establish-Timer, First-Retry-Timer, Hello-Timer, L2TP-Mode, Retry-Count

Retry-Count

Description: Specifies the maximum number of times that the MAX TNT unit attempts to establish a tunnel. Any change you make to this parameter takes effect when the previous timer expires.

Usage: Specify a decimal number from 1 to 10. The default is 10.

Example: `set retry-count = 10`

Dependencies: Retry-Count applies only if you have set L2TP-Mode to LAC. The MAX TNT can operate as a LAC only.

Location: L2-Tunnel-Global

See Also: Control-Connect-Establish-Timer, First-Retry-Timer, Hello-Timer, L2TP Mode, LAC-Incoming-Call-Timer

L2TP list attempts

With MAX TNT TAOS 8.0.0, a MAX TNT unit operating as an L2TP Network Server (LNS) can take advantage of the existing DNS List feature to attempt to connect to a series of server endpoints if the first attempt fails.

To use this feature, the MAX TNT must be configured for DNS List and the DNS servers at your site must support a list feature that enables them to return multiple addresses for a

hostname in response to a DNS query. For details about configuring DNS List, see the *MAX TNT Network Configuration Guide*.

The following example shows how to enable DNS List with a maximum of 3 hosts in the list:

```
admin> read ip-global
IP-GLOBAL read

admin> set dns-list-attempt = yes

admin> set dns-list-size = 3

admin> write
IP-GLOBAL written
```

For the MAX TNT to use DNS list when attempting to bring up a tunnel, the client's Connection or RADIUS profile must specify a DNS-resolvable hostname as the tunnel endpoint. For example, the following command shows a hostname specified as the primary tunnel-server in a Connection profile:

```
admin> get connection client-1 tunnel-options primary-tunnel-server
[in CONNECTION/client-1:tunnel-options:primary-tunnel-server]
primary-tunnel-server = tunnel-endpoint-1
```

Following is a RADIUS profile with a comparable setting:

```
5551000 Password = "Ascend-CLID", Service-Type = Dialout-Framed-User
      Tunnel-Type = L2TP,
      Tunnel-Medium-Type = IP,
      Tunnel-Server-Endpoint = "tunnel-endpoint-1"
```

When the client dials in, the system sends a DNS query to resolve the tunnel-server hostname. If it receives a list of IP addresses in return, the MAX TNT first tries to connect to the first IP address in the list. If that attempt fails, the unit continues to attempt to connect to the IP addresses in the list until a tunnel is successfully established, the DNS list has no more IP addresses, or the connection times out.

Multiple endpoints for tunnel sessions (RADIUS only)

With MAX TNT TAOS 8.0.0, when RADIUS authentication is in use, a user's profile can be configured for more than two tunnel endpoints. Each endpoint can specify its own set of attributes, such as the tunneling protocol and password. For details about this type of RADIUS configuration for tunnel connections, see "RADIUS: Tunnel attribute sets with tags and preferences" on page 70.

Secondary tunnel server for L2TP and L2F tunnels (local profiles)

With MAX TNT TAOS 8.0.0, you can configure local Connection profiles with a secondary tunnel endpoint for L2TP or L2F tunnel sessions. For details about RADIUS support for multiple endpoints, see "RADIUS: Tunnel attribute sets with tags and preferences" on page 70.

Following are the relevant parameters, shown with sample settings:

```
[in IP-GLOBAL]
system-ip-addr = 1.1.1.1

[in CONNECTION/test:tunnel-options]
profile-type = mobile-client
tunneling-protocol = l2tp-protocol
```

```
primary-tunnel-server = 2.2.2.2
secondary-tunnel-server = 3.3.3.3
```


Parameter	Specifies
System-IP-Addr	Source address for packets generated by the system. You must set the System-IP-Addr parameter in a MAX TNT unit that is operating as a LAC, particularly if the unit has multiple interfaces into the IP cloud that separates it from LNS systems. For more detail about this parameter, see the <i>MAX TNT Reference Guide</i> .
Profile-Type	For the tunnel-server parameters (next) to apply, this parameter must be set to <code>mobile-client</code> .
Tunneling-Protocol	Protocol used for tunneling. In this release, secondary tunnel server specifications are new for the <code>l2tp-protocol</code> and <code>l2f-protocol</code> settings
Primary-Tunnel-Server	IP address or hostname of the primary endpoint for L2TP or L2F tunnels. The secondary endpoint is used only if the primary server is unavailable.
Secondary-Tunnel-Server	IP address or hostname of the secondary tunnel endpoint. In previous releases, this setting was supported only for ATMP tunnels. It now applies to L2TP and L2F tunnels also.

The MAX TNT opens a tunnel session with this server only if the primary server is unavailable. Once it has established a tunnel to the secondary tunnel server, the unit maintains that tunnel until the connection terminates, even if the primary server becomes available.

Example of configuring an L2TP tunnel to two server endpoints

Figure 46 shows a MAX TNT unit that can connect to one of two possible LNS endpoints to create an L2TP tunnel for the dial-in client. In this example, the LNS endpoints are on remote networks, so the system requires Connection or RADIUS profiles to establish a connection to the endpoint systems.

Figure 46. Primary and secondary L2TP tunnel endpoints



The following commands configure the MAX TNT system IP address:

```
admin> read ip-global
IP-GLOBAL read
admin> set system-ip-addr = 1.1.1.1
```

```
admin> write
IP-GLOBAL written
```

The following commands configure Connection profiles to the two LNS systems:

```
admin> read connection endpoint-1
CONNECTION/endpoint-1 read

admin> set active = yes

admin> set dial-number = 9-1-333-555-1212

admin> set ppp-options send-password = lns-pw

admin> set ppp-options recv-password = lac-pw

admin> set ip-options remote = 2.2.2.2

admin> write
CONNECTION/endpoint-1 written

admin> read connection endpoint-2
CONNECTION/endpoint-2 read

admin> set active = yes

admin> set dial-number = 9-1-123-555-1234

admin> set ppp-options send-password = lns-pw

admin> set ppp-options recv-password = lac-pw

admin> set ip-options remote = 3.3.3.3

admin> write
CONNECTION/endpoint-2 written
```

The following commands create a Connection profile for the dial-in client:

```
admin> new connection dialin-1
CONNECTION/dialin-1 read

admin> set active = yes

admin> set clid = 555-1000

admin> set tunnel-options profile-type = mobile-client

admin> set tunnel-options tunneling-protocol = l2tp

admin> set tunnel-options primary-tunnel-server = 2.2.2.2

admin> set tunnel-options secondary-tunnel-server = 3.3.3.3

admin> write
CONNECTION/dialin-1 written
```

Parameter reference entry

Secondary-Tunnel-Server

Description: Specifies the IP address or hostname of an ATMP secondary Home Agent, or of a secondary tunnel server for an L2F or L2TP tunnel. The MAX TNT unit initiates a connection to the specified host only if the primary server is unavailable.

Usage: Specify an IP address, in dotted decimal notation, or a symbolic hostname of up to 31 characters. The default is 0.0.0.0.

If you specify a hostname, the MAX TNT uses the Domain Name System (DNS) to look up the host IP address.

For ATMP tunnels only, you can append a UDP port number to the address or hostname if the Home Agent requires the use of a port other than the one specified in the UDP-Port parameter setting. For ATMP, the port number is separated from the address or hostname by a colon. For example, 1.1.1.1:512. Do not append a port number to the value for layer 2 tunnels.

Dependencies: You must set Profile-Type to `mobile-client` for this setting to apply.

Location: Connection > Tunnel-Options

See Also: Home-Network-Name, Max-Tunnels, Password, Primary-Tunnel-Server, Profile-Type, UDP-Port

Optional L2TP system name

MAX TNT units can now use an optional system name in establishing L2TP tunnels. Following is the relevant parameter, shown with its default value:

```
[in L2-TUNNEL-GLOBAL]
l2tp-system-name = ""
```

If an L2TP system name is defined, it is used instead of the MAX TNT system name and domain name in establishing the session. If the L2TP system name is not defined, the LAC passes its actual system and domain name, as in previous releases.

The following commands set the L2TP system name:

```
admin> read l2-tunnel-global
L2-TUNNEL-GLOBAL read

admin> set l2tp-system-name = maxtnt-1

admin> write
L2-TUNNEL-GLOBAL written
```

L2TP-System-Name

Description: Specifies a name (up to 31 characters) to be passed to the LNS when initiating an L2TP tunnel. With the default null value, the LAC system name and domain name are sent.

Usage: Enter a string of up to 31 characters.

Example: `set l2tp-system-name = maxtnt-1`

Dependencies: If the name is longer than 31 alphanumeric characters, the hostname that is passed to the L2TP endpoint is truncated and ends with a plus sign (+).

Location: L2-Tunnel-Global


Limited support for Layer 2 Forwarding (L2F)

Note: This implementation of Layer 2 Forwarding (L2F) was designed to interoperate with IOS 11.3 from Cisco Systems. Other software versions or tunnel peers may not be supported.

With MAX TNT TAOS 8.0.0, a MAX TNTunit can operate as an L2F Network Access Server (NAS) in communication with an L2F home gateway that is a Cisco router running IOS 11.3.

Figure 47 shows the elements of an L2F tunnel. A PPP client dials in across an asynchronous or synchronous link, using any protocol that can be carried within PPP. The MAX TNT answers the call and passes it to the home gateway (a Cisco router running IOS 11.3). Communication between the NAS and the home gateway requires IP connectivity.

Figure 47. L2F tunneling



The connection to the home gateway is an IP link, which consists of a control link and one or more data links. The control and data links both use UDP port 1701 and are encapsulated in UDP.

The control link carries information used to query whether the home gateway can accept the current call, and to establish a tunnel. L2F implements a periodic Hello mechanism by which the NAS and home gateway verify that the other is still operational. If the Hello message doesn't arrive within a specified period, the tunnels are brought down.

Each tunneled client connection has one data link, which consists of PPP frames.

Authenticating tunnels

The MAX TNT supports both shared-secret and distinct-secret L2F tunnel authentication. The default method is to use a shared secret between the tunnel endpoints. (Authenticating L2F tunnels with a shared secret is similar to authenticating L2TP tunnels with a shared secret. For details, refer to the *MAX TNT Network Configuration Guide*.)

Distinct secrets enable you to specify different passwords to authenticate the NAS to the home gateway, and the home gateway to the NAS.

You can also configure the MAX TNT to authenticate tunnels by first attempting to use a shared secret. If that fails, the MAX TNT then authenticates tunnels using distinct secrets.

The following sequence of events describes how the MAX TNT uses distinct tunnel secrets to authenticate L2F tunnels:

- 1 A client connects and is partly authenticated. The MAX TNT looks up the associated Connection profile (or RADIUS profile) based on the client name and partially authenticates the client based on the username and password it provides.
- 2 If an L2F tunnel is specified in either the Tunnel-Type attribute in a RADIUS profile or the Tunnel-Protocol parameter in a local Connection profile, the MAX TNT either adds this client connection to an existing tunnel, or creates a new tunnel to the specified server endpoint.
- 3 If a password is present in either the Tunnel-Password attribute in RADIUS or the Password parameter in a local Connection profile, the MAX TNT uses this password to authenticate the NAS to the home gateway.

- 4 The MAX TNT authenticates the home gateway by looking up a profile that matches the name provided by the home gateway in either the Tunnel-Server-Endpoint attribute in RADIUS or the Server-Endpoint parameter in a local Connection profile.
- 5 The MAX TNT establishes the tunnel between itself and the home gateway.

If you are using RADIUS to authenticate L2F tunnels with distinct passwords, make sure of the following:

- The client's RADIUS user profile must contain a Tunnel-Password attribute with the password that the MAX TNT uses to authenticate the tunnel to the home gateway.
- The home gateway must have a RADIUS user profile. Because this is not a user profile for interactive access, Lucent recommends that the Service-Type be set to Outbound.

The following examples show a client's RADIUS profile and a home gateway's RADIUS profile that use for distinct secrets for tunnel authentication:

```
dialup-client Password = "client-pw"
    Tunnel-Type = L2F,
    Tunnel-Server-Endpoint = "1.1.1.1",
    Tunnel-Password = "nas-secret"

hg-name Password = "hg-secret", Service-Type = Outbound
    Reply-Message = ""
```

Alternatively, the home gateway password can be configured locally in a Tunnel-Server profile.

Configuring basic L2F operations

To enable the MAX TNT to operate as an L2F endpoint, you must set it to run in NAS mode and configure it to recognize the L2F home gateway (a Cisco router running IOS 11.3).

If the home gateway is on a remote IP network, the MAX TNT also requires an IP-routed Connection profile or RADIUS profile to the home gateway. For details about configuring IP WAN interfaces, see the *MAX TNT Network Configuration Guide*.

Overview of global L2F parameters

Following are the global L2F parameters (shown with default values) for configuring L2F operations:

```
[in L2-TUNNEL-GLOBAL]
udp-queue-length = 256
l2f-mode = disabled
l2f-system-name = ""
l2f-retry-count = 4
l2f-retry-interval = 0
l2f-tunnel-secret = ""

[in TUNNEL-SERVER/" "]
server-endpoint* = ""
enabled = yes
shared-secret = ""
```

Parameter	Specifies
UDP-Queue-Length	Maximum number of UDP packets that can reside in the input queue for the L2F NAS. The default queue length for UDP requests is 256. Valid values for the queue length are 0–512.
L2F-Mode	Enable/disable L2F operations. Specify NAS to enable L2F operations in the MAX TNT. The default is Disabled.
L2F-System-Name	System name of the NAS unit. It is used to identify the NAS to the L2F home gateway during tunnel creation.
L2F-Retry-Count	Number of times the MAX TNT will resend L2F control packets. Values can be from 1 to 16. The default is 4.
L2F-Retry-Interval	Retry interval in seconds. Values can be from 0 to 32 seconds. The default value of 0 specifies that an adaptive retry interval, (based on the retry number plus 1) is to be used.
L2F-Tunnel-Secret	Authentication method used by the MAX TNT to authenticate the L2F tunnels. When the parameter is set to <code>shared-tunnel-secret</code> (the default), tunnel authentication relies on a secret shared by the NAS and the home gateway. When set to <code>distinct-tunnel-secrets</code> , tunnel authentication uses distinct secrets for authenticating the NAS to the home gateway, and the home gateway to the NAS. When set to <code>either-shared-or-distinct-tunnel-secret</code> , the MAX TNT first tries to authenticate using the shared secret. If that fails, the unit then tries to authenticate the tunnel using distinct secrets. For more information, see “Authenticating tunnels” on page 294.
Server-Endpoint	DNS hostname or dotted-decimal IP address of the tunnel endpoint. If it specifies a hostname, the MAX TNT executes a DNS lookup for the host’s address.
Enabled	Enable/disable tunnels to the specified Server-Endpoint.
Shared-Secret	Shared secret for authenticating L2F tunnels. L2F tunnels can be authenticated with the same secret value at both ends of the connection.

Example of a basic L2F configuration

The following commands configure basic L2F operations with a home gateway named l2f-1:

```
admin> read l2-tunnel
L2-TUNNEL-GLOBAL read
admin> set l2f-mode = NAS
admin> set l2f-system-name = l2f-1
admin> write
L2-TUNNEL-GLOBAL written
admin> read tunnel-server l2f-1
TUNNEL-SERVER/l2f-1 read
admin> set server-endpoint = 1.1.1.1
admin> set enabled = yes
admin> set shared-secret = secret1
```



```
admin> write
TUNNEL-SERVER/l2f-1 written
```

Configuring L2F client profiles

When a PPP client dials into the MAX TNT to initiate a tunnel to the L2F home gateway, the MAX TNT must first authenticate the client by PPP authentication. Even though the MAX TNT has provided password authentication for a call, the home gateway can (and probably should, for security reasons) authenticate again. The NAS and home gateway can use different PPP authentication protocols without restriction.

Overview of L2F settings in Connection profiles

Following are the L2F tunnel parameters (shown with sample values) in a Connection profile:

```
[in CONNECTION/tunnelcx:tunnel-options]
profile-type = mobile-client
tunneling-protocol = l2f-protocol
primary-tunnel-server = l2f-1
password = "nas-pw"
```

Parameter	Specifies
Profile-Type	Type of tunnel profile. Specify Mobile-Client for L2F tunneling.
Tunneling-Protocol	Protocol to use when creating a tunnel for this profile. Set to <code>l2f-protocol</code> to pass traffic to a home gateway.
Primary-Tunnel-Server	DNS hostname or dotted-decimal IP address of the home gateway endpoint. If this parameter specifies a hostname, the MAX TNT executes a DNS lookup for the host's address. If DNS List is supported and this parameter specifies a hostname, the MAX TNT continues to attempt to connect to the IP addresses in the list until a tunnel is successfully established, the DNS list has no more IP addresses, or the connection times out. For more details, see "L2TP List Attempts" on page 236.
Password	Shared secret for authenticating L2F tunnels. If the Password parameter is configured, the MAX TNT uses it to authenticate L2F tunnels and ignores the Shared-Secret setting in the Tunnel-Server profile. The Password is used only for establishing the tunnel and is ignored for subsequent connections added to the tunnel.

Overview of L2F settings in RADIUS profiles

RADIUS uses the following attribute-value pairs to specify L2F tunnels:

RADIUS attribute	Value
Tunnel-Type (64)	Tunneling protocol to be used. Set to L2F (2) for L2F tunneling.
Tunnel-Medium-Type (65)	Medium to be used for the tunnel. Only IP (1) is supported at this time.

RADIUS attribute	Value
Tunnel-Server-Endpoint (66)	DNS hostname or dotted IP address of the home gateway endpoint (a string value). If this attribute specifies a hostname, the MAX TNT executes a DNS lookup for the host's address. If DNS List is supported and this attribute specifies a hostname, the MAX TNT continues to attempt to connect to the IP addresses in the list until a tunnel is successfully established, the DNS list has no more IP addresses, or the connection times out. For more details, see "L2TP List Attempts" on page 236.
Tunnel-Password (69)	Shared secret for authenticating L2F tunnels. If distinct secrets are being used, Tunnel-Password is the secret used by the NAS to authenticate to the home gateway.
Tunnel-Client-Auth-ID (90)	Name to be used as the NAS name during L2F tunnel establishment. Note that this can be different than the L2F-System-Name. If configured, Tunnel-Client-Auth-ID overrides the L2F-System-Name parameter.

Examples of opening a tunnel after password authentication

In this example, the MAX TNT negotiates the PPP call, including password authentication, and then opens the L2F tunnel. For details about PPP authentication, see the documentation that came with your unit.

The following commands create a Connection profile that includes a PPP password. The MAX TNT authenticates the caller before bringing up a tunnel to a home gateway at 1.1.1.1:

```
admin> read conn l2test
CONNECTION/l2test read
admin> set active = yes
admin> set encapsulation-protocol = ppp
admin> set ppp-options rcv-password = localpw
admin> set tunnel-options profile-type = mobile-client
admin> set tunnel-options primary-tunnel-sever = 1.1.1.1
admin> set tunnel-options tunneling-protocol = l2f-protocol
admin> write
CONNECTION/l2test written
```

Following is a comparable RADIUS profile:

```
l2test Password = "localpw"
  Service-Type = Framed,
  Framed-Protocol = PPP,
  Tunnel-Server-Endpoint = "1.1.1.1",
  Tunnel-Type = L2F,
  Tunnel-Medium-Type = IP,
  Tunnel-Password = "shared_secret"
```

Parameter reference entries

The following existing parameters (shown with default values) now also relate to L2F tunnels:

```
[in TUNNEL-SERVER/""]
server-endpoint* = ""
enabled = yes
shared-secret = ""

[in CONNECTION/"":tunnel-options]
profile-type = disabled
tunneling-protocol = atmp-protocol
primary-tunnel-server = ""
secondary-tunnel-server = ""
password = ""
```

In addition, the following parameters have been added to the MAX TNT interface to support L2F.

L2F-Mode

Description: Enables or disables L2F operations. Currently, the MAX TNT can only operate in NAS mode.

Usage: Specify one of the following values:

- NAS enables L2F on the MAX TNT.
- Disabled (the default) disables L2F on the MAX TNT.

Location: L2-Tunnel-Global profile

See Also: L2F-System-Name, L2F-Retry-Count, L2F-Retry-Interval, Tunnel-Secret, UDP-Queue-Length

L2F-Retry-Count

Description: The number of times the MAX TNT resends L2F control packets.

Usage: Specify a number from 1 to 16. The default is 4.

Example: `set l2f-retry-count = 8`

Dependencies: L2F-Retry-Count applies only if L2F-Mode is set to nas.

Location: L2-Tunnel-Global profile

See Also: L2F-Mode, L2F-System-Name, L2F-Retry-Interval, Tunnel-Secret, UDP-Queue-Length

L2F-Retry-Interval

Description: The retry interval in seconds. Values can be from 0 to 32 seconds.

Usage: Specify a number between 0 and 32. The default value of 0 specifies that an adaptive retry interval, (based on the retry number plus 1) is to be used.

Example: `set l2f-retry-interval = 4`

Dependencies: L2F-Retry-Interval applies only if L2F-Mode is set to nas.

Location: L2-Tunnel-Global profile

See Also: L2F-Mode, L2F-System-Name, L2F-Retry-Count, Tunnel-Secret, UDP-Queue-Length

L2F-System-Name

Description: System name of the NAS unit. It is used to identify the NAS to the L2F home gateway during tunnel creation.

Usage: Specify a system name of up to 24 characters.

Dependencies: L2F-System-Name only applies if L2F-Mode is set to nas.

Example: `set l2f-system-name = l2f-ny-1`

Location: L2-Tunnel-Global profile

See Also: L2F-Mode, L2F-Retry-Count, L2F-Retry-Interval, Tunnel-Secret, UDP-Queue-Length

Tunnel-Secret

Description: The authentication method used by the MAX TNT to authenticate the L2F tunnels.

Usage: Specify one of the following values:

- Shared-Tunnel-Secret (the default): Tunnel authentication relies on a secret shared by the NAS and the home gateway.
- Distinct-Tunnel-Secrets: Tunnel authentication uses distinct secrets for authenticating the NAS to the home gateway, and the home gateway to the NAS.
- Either-Shared-or-Distinct-Tunnel-Secret: The MAX TNT first tries to authenticate using the shared secret. If that attempt fails, the unit then tries to authenticate the tunnel using distinct secrets.

Example: `set tunnel-secret = distinct-tunnel-secrets`

Dependencies: Tunnel-Secret applies only if L2F-Mode is set to nas.

Location: L2-Tunnel-Global profile

See Also: L2F-Mode, L2F-System-Name, L2F-Retry-Count, L2F-Retry-Interval, UDP-Queue-Length

UDP-Queue-Length

Description: The maximum number of UDP packets that can reside in the input queue for the L2F NAS.

Usage: Specify a value from 0 to 512. The default is 256. Zero (0) specifies that the packets are not dropped, no matter how busy the UDP subsystem gets. Use the zero value with caution, however. If the queue grows too large in an extremely loaded routing environment, the system can ultimately run out of memory.

Example: `set udp-queue-length = 512`

Location: L2-Tunnel-Global profile

See Also: L2F-Mode, L2F-System-Name

Virtual routing

Virtual router (VRouter) DNS

With MAX TNT TAOS 8.0.0, you can configure and manage DNS information separately for each virtual router (VRouter), to completely segment the VRouter's DNS information from any other hosts. VRouter DNS configuration includes settings for primary and secondary DNS servers, domain names, and client DNS servers, for directing connections that belong to the VRouter to a particular DNS service. The addresses configured for client DNS servers are presented to dial-in users during IP Control Protocol (IPCP) negotiation.

If DNS information is not found in the VRouter profile, the system uses the DNS information in the IP-Global profile, as in previous releases. The DNS list and the local DNS table maintained in RAM are system-wide DNS configurations that are not supported separately for each VRouter.

Overview of VRouter DNS settings

Following are the VRouter-specific DNS parameters, shown with their default settings:

```
[in VROUTER/" "]
domain-name = " "
sec-domain-name = " "
dns-primary-server = 0.0.0.0
dns-secondary-server = 0.0.0.0
client-primary-dns-server = 0.0.0.0
client-secondary-dns-server = 0.0.0.0
allow-as-client-dns-info = True
```

Parameter	Specifies
Domain-Name	Primary domain name to use for DNS lookups for this VRouter, up to 63 characters. The MAX TNT appends this domain name to hostnames when performing lookups.
Sec-Domain-Name	Secondary domain name to use for DNS lookups for this VRouter if the hostname is not found in the primary domain.
DNS-Primary-Server	Address of the primary local DNS server to use for lookups for this VRouter.
DNS-Secondary-Server	Address of the secondary local DNS server to use for lookups for this VRouter. Used only if the primary server is not found.
Client-DNS-Primary-Server	Address of a client DNS server for dial-in clients of this VRouter.
Client-DNS-Secondary-Server	Address of a secondary DNS server for dial-in clients of this VRouter.
Allow-As-Client-DNS-Info	Enable/disable use of main (local) DNS information if the client DNS servers are not found. To isolate local network information for this VRouter, set to <code>false</code> .

Example of a typical VRouter DNS configuration

For example, the following commands specify a primary and secondary domain name for DNS lookups for a VRouter named xyz:

```
admin> read vrouter xyz
VROUTER/xyz read

admin> set domain-name = xyz.com
admin> set sec-domain-name = eng.xyz.com
admin> write
VROUTER/xyz written
```

If a lookup fails in the first domain, the router tries again with the secondary domain name. To enable the MAX TNT to look up addresses via DNS, specify DNS server addresses as shown in the following example:

```
admin> read vrouter xyz
VROUTER/xyz read

admin> set dns-primary-server = 1.2.2.2
admin> set dns-secondary-server = 1.3.3.3
admin> write
VROUTER/xyz written
```

If the primary server is unavailable, the MAX TNT attempts a lookup on the secondary server. The following commands configure a client DNS server for this VRouter:

```
admin> read vrouter xyz
VROUTER/xyz read

admin> set client-dns-primary-server = 1.2.2.2
admin> set client-dns-secondary-server = 1.2.2.96
admin> set allow-as-client-dns-info = false
admin> write
VROUTER/xyz written
```

The secondary server is accessed only if the primary one is inaccessible. If both of these client DNS servers are not accessible, the MAX TNT does not allow the client to access local DNS servers.

Modified administrative commands for VRouter DNS

The NSlookup, Ping, and ARPTable commands can now take a VRouter argument and resolve names using that VRouter's DNS information. Following are revised usage statements for the commands:

```
Usage : nslookup [-r vRouterName] hostname
```

```
Usage: ping [-qvfl] [-c count] [-i sec | -I msec] [-s packet-size] [-r vRouter] [-x source_address] host-name
```

```
Usage: arptable [vRouter] [[-a hostname MAC_address] | [-d hostname] | [-f]]
```


If the VRouter name is not specified, the DNS configured for the global (main) VRouter is used.

VRouter support for L2TP connections

In previous releases, L2TP tunnels always used the main VRouter only. With MAX TNT TAOS 8.0.0, L2TP tunnels can now be built on specific VRouters. L2TP packets (control channel and encapsulated data) are sent using the configured VRouter for that tunnel.

Because each VRouter maintains its own routing table and knows about only those interfaces that explicitly specify the same VRouter, this feature allows the system to separate traffic for different LNS systems. For example, Figure 48 shows two dial-in clients, MC-1 and MC-2. Each client tunnels to a different LNS, but both LNS systems have the IP address 1.1.1.1. Because the tunnels are built on separate VRouters, the traffic is kept separate and directed to the appropriate server endpoint.

Figure 48. L2TP tunnels built on separate VRouters



Note that MAX TNT must dedicate one IP interface to each VRouter. Following are the parameters, shown with sample values, for dedicating an Ethernet or a WAN IP interface to a VRouter:

```
[in IP-INTERFACE/{ { shelf-1 slot-3 1 } 0 }]
vrouter = VRouter-1

[in CONNECTION/LNS-2]
vrouter = VRouter-2
```

For details about L2TP and VRouters, see the *MAX TNT Network Configuration Guide*.

Connection profile setting

Following is the parameter (shown with its default value) for specifying a VRouter name:

```
[in CONNECTION/MC-1:tunnel-options]
vrouter = ""
```

Parameter	Specifies
VRouter	Name of a virtual router to use for establishing the L2TP tunnel. The specified VRouter must exist on the LAC. With the default null value, the global VRouter is used.

For example, the following commands configure a mobile-client profile for an L2TP session that belongs to a VRouter named VRouter-1:

```
admin> new connection MC-1
CONNECTION/MC-1 read
admin> set active = yes
admin> set encapsulation-protocol = ppp
```

Extension features in MAX TNT TAOS 8.0.0

Short-duration transaction network (SDTN)

```
admin> set ppp-options recv-password = localpw
admin> set tunnel-options profile-type = mobile-client
admin> set tunnel-options primary-tunnel-server = 1.1.1.1
admin> set tunnel-options tunneling-protocol = l2tp
admin> set tunnel-options vrouter = VRouter-1
admin> write
CONNECTION/MC-1 written
```

With this sample profile, the MAX TNT authenticates the caller before building a tunnel to the LNS at 1.1.1.1 on the specified VRouter.

RADIUS profile setting

RADIUS uses the following attribute-value pair to specify a VRouter name:

RADIUS attribute	Value
Ascend-Tunnel-VRouter-Name (31)	Name of a virtual router to use for establishing the L2TP or L2F tunnel. The specified VRouter must exist on the LAC. With the default null value, the global VRouter is used. This attribute may be part of an attribute set by including a tag. For details, see “RADIUS: Tunnel attribute sets with tags and preferences” on page 70.

For example, the following mobile-client profile specifies an L2TP session that belongs to a VRouter named VRouter-2:

```
MC-2 Password = "localpw"
Service-Type = Framed-User,
Framed-Protocol = PPP,
Tunnel-Server-Endpoint = "1.1.1.1",
Tunnel-Type = L2TP,
Ascend-Tunnel-VRouter-Name = "VRouter-2"
```

Short-duration transaction network (SDTN)


With MAX TNT TAOS 8.0.0, the MAX TNT unit supports interaction with transaction servers to conduct short-duration transactions over IP-based networks. SDTN is a hash-code protected feature. The related parameters might be visible in the command-line interface but are not enabled unless the appropriate software license has been purchased from Lucent Technologies. You can verify that the SDTN license is enabled in your default Transaction-Server profile by using the following command:

```
admin> get transaction-server enabled
enabled = yes
```

SDTN operation

To support short-duration transactions, the MAX TNT receives calls from transaction client applications and transparently forwards them to a transaction server. Figure 49 shows a sample SDTN setup, with transaction servers on a local 100-Mb Ethernet interface.

Figure 49. Sample SDTN setup



Transaction data calls come in from National Institute of Standards and Technology (NIST) or Connectionless Network Protocol (CLNP) terminals, or Visa terminals. The MAX TNT answers the calls and forwards them to the transaction server via the Quick Transaction Protocol (QTP).

QTP is a symmetrical protocol that operates over UDP in both directions between the MAX TNT and transaction servers. QTP establishes and releases the virtual connection between systems, transports transaction traffic, and exchanges periodic Status Report messages.

To determine which server to use for a particular transaction processing request, the MAX TNT uses a selection table. The system keeps the table up-to-date on server availability and status by applying configurable metrics to information obtained from QTP Status Report messages and from real-time events, such as failure to receive a response to a call request.

Transaction-Server profiles

The Transaction-Server profile sets parameters that affect the metrics used in the server selection table. The table contains a primary and secondary list of transaction servers that have been entered into the list via QTP. The MAX TNT uses only the primary list unless no available servers are left in the primary list, in which case it begins using the secondary list. When a server adds itself to the list, the MAX TNT generates one of the following Syslog messages:

```
TS Address [x.x.x.x] has been entered into the Primary List
TS Address [x.x.x.x] has been entered into the Secondary List
```

Each list entry specifies a transaction server's IP address, the UDP port used by QTP on that server, and a metric that indicates the server's availability to the MAX TNT. In this release, the MAX TNT searches the list in cyclic order and chooses the first available server. (The metric is not used to weight the selection in this release. It is used to remove servers from the list when their status or availability crosses a metric threshold. In future releases, additional hunt mechanisms will be supported.) When the MAX TNT removes itself from the list, it generates one of the following Syslog messages:

```
TS Address [x.x.x.x] has been removed from Primary List
TS Address [x.x.x.x] has been removed from Secondary List
```

The parameter settings in the Transaction-Server profile are used to associate metrics with the events that keep the table up-to-date: QTP Status Report messages, events such as call requests and responses, and periodic receipt of QTP Status Reports. If these events do not occur as expected, the system can change a transaction server metric on that basis.

Extension features in MAX TNT TAOS 8.0.0

Short-duration transaction network (SDTN)

The QTP Status Report messages from transaction servers can contain the following flow control attributes, indicating how busy the server is:

- Available (0x01)
- Partly Congested (0x02)
- Congested (0x03)
- Shutdown (0x04)

QTP Status Report messages can also contain the Primary Station (0x01) or Secondary Station (0x02) status attribute, indicating whether the server is on the primary or secondary list.

Overview of transaction server settings

The following parameters, shown here with default values, are used to configure the metric algorithms and thresholds used for transaction server selection:

```
[in TRANSACTION-SERVER]
enabled = yes
hunting-mechanism = cyclic
selection-timeout = 10000
data-ack-timeout = 10000
keep-alive-timeout = 30
qtp-port = 3350
metric-max = 15
no-conn-ack-increment = 8
call-reject-increment = 4
call-ack-decrement = 1
available-metric = 1
partly-congested-metric = 4
congested-metric = 10
shutdown-metric = 14
no-first-status-metric = 10
no-second-status-metric = 16
max-qtp-pdu-size = 512
```


Parameter	Specifies
Enabled	Status of the SDTN license (read-only). It is set to <code>yes</code> when the license is enabled. If its value is <code>No</code> , the license is disabled and this profile has no effect.
Hunting-Mechanism	The method by which to search the Primary list (or Secondary list) of transaction servers. In this release, only the default <code>cyclic</code> setting is supported, which indicates that the list is searched in cyclic order.
Selection-Timeout	Number of milliseconds (from 0 to 65000) before the attempt to establish a QTP connection with a transaction server times out. The default is 10000 milliseconds.
Data-Ack-Timeout	Number of milliseconds (from 500 to 30000) that the MAX TNT waits for a transaction server to send a QTP Acknowledge in response to a QTP data message. The default is 10000 milliseconds.
Keep-Alive-Timeout	Number of seconds (from 1 to 300) that the MAX TNT waits for a QTP Status update from a transaction server. The default is 30 seconds.

Parameter	Specifies
QTP-Port	UDP port for QTP to listen for incoming QTP connections. UDP port numbers can be from 0 to 65535. The default port number for QTP is 3350.
Metric-Max	Number from 0 to 255 indicating the maximum metric, beyond which a transaction server is removed from an active list. The default maximum metric is 15.
No-Conn-Ack-Increment	Number from 0 to 255 by which to increase a transaction server's current metric if it does not send a QTP Connect Acknowledgement in response to a QTP Connect Request sent by the MAX TNT. The default setting is 8.
Call-Reject-Increment	Number from 0 to 255 by which to increase a transaction server's current metric if it sends a QTP Call Reject in response to a QTP Connect Request sent by the MAX TNT (a QTP connection attempt failed). The default setting is 4.
Call-Ack-Decrement	Number from 0 to 255 by which to decrease a transaction server's current metric if it sends a QTP Call Ack in response to a QTP Connect Request sent by the MAX TNT (a QTP connection attempt succeeded). The default setting is 1.
Available-Metric	Number from 0 to 255 to use as a transaction server's current metric if it sends a QTP Status Message with a Flow Control Attribute set to Available. The default setting is 1.
Partly-Congested-Metric	Number from 0 to 255 to use as a transaction server's current metric if it sends a QTP Status Message with a Flow Control Attribute set to Partly-Congested. The default setting is 4.
Congested-Metric	Number from 0 to 255 to use as a transaction server's current metric if it sends a QTP Status Message with a Flow Control Attribute set to Congested. The default setting is 10.
Shutdown-Metric	Number from 0 to 255 to use as a transaction server's current metric if it sends a QTP Status Message with a Flow Control Attribute set to Shutdown. The default setting is 14.
No-First-Status-Metric	Number from 0 to 255 to use as a transaction server's current metric the first time it does not send a QTP Status Message within the timeout interval. The default setting is 10.
No-Second-Status-Metric	Number from 0 to 255 to use as a transaction server's current metric the second time it does not send a QTP Status Message within the timeout interval. The default setting is 16.
Max-QTP-PDU-Size	Maximum number of bytes (from 1 to 1460) a QTP message sent by the MAX TNT can contain. The default is 512 bytes.

Example of a transaction server configuration

Figure 50 shows a sample SDTN setup with two MAX TNT units. For redundancy purposes as well as speed of access, each MAX TNT unit and transaction server supports a 100-Mbps Ethernet interface on two local subnets. Because QTP Status Reports from the transaction servers contain the IP addresses of both Ethernet interfaces on each server, a single server appears as two addressable entities in the server selection table. These connections provides some redundancy if a failure occurs on one subnet or Ethernet port, because the server is still reachable on the other subnet or port.

Figure 50. Transaction servers with redundant Ethernet connections



For most sites, the default settings in the Transaction-Server profile are the most effective SDTN setup.

Dial-in connections for transaction clients

The MAX TNT recognizes HDLC-Normal Response Mode (HDLC-NRM) and Visa-II dial-in connections for transaction processing.

Any transaction client connection requires quick handling to avoid timeouts. If the call is made via modem, you can configure a custom AT string that specifies the required modem timings, modulation types, speed, and other modem parameters. This customization helps to prevent delays caused by modem training.

Answer-Defaults profile settings

The Answer-Defaults profile contains two new subprofiles for the HDLC-NRM and Visa-II link-layer encapsulation protocols. Following are the relevant parameters, shown with default settings:

```
[in ANSWER-DEFAULTS:hdlc-nrm-answer]
enabled = yes

[in ANSWER-DEFAULTS:visa2-answer]
enabled = yes
```

By default, the system does not reject HDLC-NRM or Visa-II calls on the basis of their encapsulation types. With the default settings, the system answers the calls if they pass authentication.

For HDLC-NRM and Visa-II connections, CLID or DNIS authentication is required. For example, the following commands configure the unit to require DNIS:

```
admin> read answer
ANSWER-DEFAULTS read

admin> set clid-auth-mode = dnis-require

admin> write
ANSWER-DEFAULTS written
```

For details about CLID and DNIS authentication, see the *MAX TNT Network Configuration Guide*.

Configuring HDLC-NRM connections

The MAX TNT now supports HDLC-NRM as a link-layer encapsulation protocol. When it receives an HDLC-NRM call, the system must first authenticate the call via CLID or DNIS. If the call passes authentication, the system answers it, completes HDLC negotiations, and forwards the packets to its QTP software, which routes the packets via UDP to a transaction server.

HDLC-NRM is similar to the Link Access Procedure, Balanced (LAPB) protocol and other layer 2 HDLC protocols. The initial HDLC-NRM packet is SNRM (Set Normal Response Mode). Unlike LAPB, in which the connected stations are peers and each has the right to send data at any time, HDLC-NRM is a half-duplex protocol, so only one station is allowed to send data at a time. To enable this, one of the connected stations is the primary station (typically the MAX TNT) and the other is the secondary station (typically the NIST or CLNP terminal). The primary station can send data packets at any time. The secondary station must be polled (via RR) before it can send data packets as synchronous I-frames. By default, the primary station drops data packets it receives from the secondary station when the station has not been given the right to transmit (asynchronous I-frames).

Overview of HDLC-NRM settings

The following parameters, shown with default values, are used to configure HDLC-NRM connections:

```
[in CONNECTION/dgtnt]
encapsulation-protocol = hdlc-nrm
sdtn-packets-server = no

[in CONNECTION/dgtnt:hdlc-nrm-options]
enabled = no
snrm-response-timeout = 20000
snrm-retry-counter = 2
poll-timeout = 60000
poll-rate = 5000
poll-retry-counter = 2
primary = yes
async-drop = yes
```

Parameter	Specifies
Encapsulation-Protocol	Encapsulation protocol to use for this connection. Must be set to <code>hdlc-nrm</code> for HDLC-NRM clients.
SDTN-Packets-Server	Enable/disable forwarding packets to a transaction server via QTP. Set this parameter to <code>yes</code> for HDLC-NRM connections. If set to <code>no</code> (the default) for an HDLC-NRM connection, the system establishes the connection but drops the data.
Enabled	Enable/disable answering of an HDLC-NRM call that matches this profile.
SNRM-Response-Timeout	SNRM (Set Normal Response Mode) is the initial HDLC-NRM packet sent. This setting specifies the number of milliseconds (500-5000) to wait for a response. The default is 2000.
SNRM-Retry-Counter	Number of times (from 0 to 255) to retry sending an SNRM packet following a response timeout. The default setting is 2.

Parameter	Specifies
Poll-Timeout	Number of milliseconds (from 0 to 255000) to wait for a response from the caller (the secondary station) to a poll sent by the MAX TNT (the primary station). The default setting is 60000.
Poll-Rate	Number of milliseconds (from 500 to 5000) between polls. The default setting is 5000.
Poll-Retry-Counter	Number of times (from 0 to 255) to retry the poll after a response timeout. The default setting is 2.
Primary	Primary or secondary station status of the dial-in unit. The default is no because dial-in terminals usually act as secondary stations. Setting this parameter to yes causes the MAX TNT to act as the secondary station for this connection (usually for test purposes).
Async-Drop	Because HDLC-NRM is a half-duplex protocol, the primary station must drop asynchronous I-frames it receives from a secondary station. When this parameter is set to yes (the default) and the MAX TNT is the primary station, the system drops I-frames received from the secondary station. If the parameter is set to no, the system processes the I-frames it receives normally. Setting the parameter to no enables back-to-back testing on the MAX TNT.

Example of a typical HDLC-NRM client configuration

For example, the following commands configure a Connection profile for an HDLC-NRM client:

```
admin> new conn hstation-1
CONNECTION/hstation-1 read
admin> set active = yes
admin> set encapsulation-protocol = hdlc-nrm
admin> set sdtn-packets-server = yes
admin> set dial-number = 853784
admin> set calledNumber = 3783
admin> set telco-options dialout-allowed = yes
admin> set hdlc-nrm-options enabled = yes
admin> write
CONNECTION/hstation-1 written
```

Configuring Visa-II connections

The MAX TNT now supports Visa-II as a link layer encapsulation protocol. When it receives a call from a Visa terminal, the system must first authenticate the call via CLID or DNIS. If the call passes authentication, the system answers it and forwards the packets to its QTP software, which routes the packets via UDP to a transaction server.

For Visa-II connections, protocol handling occurs between the transaction server and the Visa terminal. For incoming data from the terminal, the MAX TNT performs some minimal parsing as defined by the Visa-II settings in the caller's Connection profile. For data from the server to the terminal, the MAX TNT simply passes the data transparently.

Overview of Visa-II settings

The following parameters, shown with default values, are used to configure Visa-II:

```
[in CONNECTION/dgtnt]
encapsulation-protocol = visa2
sdtn-packets-server = no

[in CONNECTION/dgtnt:visa2-options]
enabled = no
idle-character-delay = 10000
first-data-forward-character = 04
second-data-forward-character = 06
third-data-forward-character = 15
fourth-data-forward-character = 05
1-char-sequence = 03
2-char-sequence = 00:03:00:00
```

Parameter	Specifies
Encapsulation-Protocol	Encapsulation protocol to use for this connection. Must be set to <i>visa2</i> for Visa terminal connections.
SDTN-Packets-Server	Enable/disable forwarding packets to a transaction server via QTP. Set this parameter to <i>yes</i> for Visa terminal connections. If set to <i>no</i> (the default) for a Visa terminal connection, the system establishes the connection but drops the data.
Enabled	Enable/disable answering of an Visa-II call that matches this profile.
Idle-Character-Delay	Number of milliseconds of idle time to wait after receiving a character before forwarding data. Range 0 to 30,000 ms. The default setting is 10,000 ms.
First-Data-Forward-Character	Hexadecimal value of a character to be used as a trigger to forward data. The default setting is 04.
Second-Data-Forward-Character	Hexadecimal value of a character to be used as a trigger to forward data. The default setting is 06.
Third-Data-Forward-Character	Hexadecimal value of a character to be used as a trigger to forward data. The default setting is 15.
Fourth-Data-Forward-Character	Hexadecimal value of a character to be used as a trigger to forward data. The default setting is 05.
1-Char-Sequence	Hexadecimal value of a character to be used as a trigger to forward data and the next following character. The default setting is 03.
2-Char-Sequence	Two character values of a sequence to be used as a trigger to forward data and the two following characters after the sequence. The default setting is 00:03:00:00. Note that only the first two character values in this sequence have meaning. The last two values are ignored.

Example of a typical Visa-II terminal configuration

For example, the following commands configure a Connection profile for a Visa terminal or terminal-emulator:

Extension features in MAX TNT TAOS 8.0.0

Short-duration transaction network (SDTN)

```
admin> new conn visa-1
CONNECTION/visa-1 read

admin> set active = yes

admin> set encapsulation-protocol = visa2

admin> set sdtm-packets-server = yes

admin> set dial-number = 853784

admin> set calledNumber = 34343

admin> set telco dialout-allowed = yes

admin> set visa2 enabled = yes

admin> write
CONNECTION/visa-1 written
```

Preventing training delays for modem transaction calls

When a transaction call is initiated or answered by a modem, the MAX TNT must train the modem before establishing the connection. To enable dial-in terminals for transaction processing to connect quickly with as little modem training as possible, you can specify an AT string that specifies the required modem timings, modulation types, speed, and other modem parameters. This customization helps to prevent delays caused by modem training.

Setting for customizing the AT string

Following is the relevant parameter, shown with its default value:

```
[in CONNECTION/" "]
at-string = " "
```

With the default null value, the system performs modem training as usual. The value of this parameter can be set to valid AT commands of up to 58 characters. Do not begin the string with AT. An AT is automatically appended to the beginning of this string before it is sent to the modem. Also, do not include an A (answer) or a D (dial) command anywhere in the string. An A command is automatically appended to the end of this string for incoming calls, and a D command in the answer string will cause the call to fail. A D command is appended to the end of the specified string for outgoing calls.

Note: Be very careful when entering AT commands in this parameter. The system does not prevent you from entering incorrect strings.

Example customized AT string

The following commands configure an HDLC-NRM connection and set the AT-String to force the modem to answer as a Bell 212A type modem:

```
admin> new conn hstation-1
CONNECTION/hstation-1 read

admin> set active = yes

admin> set encapsulation-protocol = hdlc-nrm

admin> set dial-number = 853784

admin> set calledNumber = 3783

admin> set telco dialout-allowed = yes
```

```
admin> set hdlc enabled = yes
admin> set at-string = B1+MS=69,1,1200,1200;
admin> write
CONNECTION/hstation-1 written
```

The sample AT-String setting causes the following string to be sent to the modem, which forces it to answer as Bell 212A type modem in automode.

```
ATB1+MS=69,1,1200,1200;
```

Extension features in MAX TNT TAOS 8.0.0

Short-duration transaction network (SDTN)

Corrections in MAX TNT TAOS 8.0.0

TR	Problem corrected
1612	Log window displayed extra lines for Windows95 Telnet sessions when the user pressed the up-arrow key to scroll through the log status.
2180	Cutting and pasting the argument to the Open command resulted in a double carriage return, which caused the system to generate an "Error from far end" message.
2490	The Quit command was included as the last entry in a configuration file saved either by using the Save command in the CLI or from an SNMP management station.
2255	In releases prior to 7.0.3, multishelf systems maintained multiple dates and times. Now, the slave systems are synchronized with the master shelf clock whenever a slave comes up, a user changes the Timedate profile on the master shelf, or the master receives SNTP updates.
3003	In previous releases, administrators could not log into a user's unit for configuration purposes. This has been fixed.
3045	SNMP: ifAdminStatus returned Up for a disabled T1 interface.
3207	When Pool-Summary is enabled and the IP address pools are defined in the IP-Global profile, the MAX TNT did not add summary routes to the routing table, so the routes were not advertised in RIP updates.
3553	The MAX TNT incorrectly forwarded broadcast packets received from a WAN link to the shelf's Ethernet network.
3626	When using TFTP to load or save a large number of profiles, the TFTP operation could sometimes time out and abort.
3739	When using TFTP to load or save a large number of profiles, the TFTP operation could sometimes time out and abort.
3746	System did not use multipath routing to load balance between interfaces.
3976	When a power supply failed or was unplugged, the MAX TNT generated a power supply state change trap (Ascend enterprise trap 24), but did not generate the trap when power was restored.
4003	MAX TNT units did not reestablish a nailed connection correctly after it had been dropped.
4043	The MAX TNT unit stopped taking fax calls after 16 hours, generating Warning messages 179 and 154 in Diagnostic mode.
4148	MAX TNT unit sent RIP packets with a Time To Live (TTL) value of 64 instead of a value of 1.
4179	MAX TNT unit did not propagate changes to a filter applied to an Ethernet interface to the interface when the Filter profile was written.
4189	The NAS-Port attribute-value pair was missing in Accounting records for calls received on shelf 1, slot 1, port 5, channel 1.

Corrections in MAX TNT TAOS 8.0.0

TR	Problem corrected
4268	MAX TNT RADIUS Statistics were not aggregated for all host slot cards.
4270	Occasionally, the MAX TNT reported progress code 65 (LCP_OPENED) instead of progress code 60 (LAN_SESS_UP) when an NCP connection was established during a PPP dial-in.
4284	MAX TNT unit stopped accepting new calls and sent busy signals.
4294	MAX TNT units supporting LAN modems did not send ASCEND_CD ON/OFF notifications to notify remote clients of carrier detect changes.
4341	Occasionally, a MAX TNT supporting Series56 II modem cards could reset, generating a Fatal Error 1 in Diagnostic mode.
4357	MAX TNT units supporting LAN modems did not send ASCEND_CD ON/OFF notifications to notify remote clients of carrier detect changes.
4358	The MAX TNT did not generate power supply traps following a system reset unless you set the powerSupplyStateTrapState and powerSupplyOperationalStateTrapState variables in the power supply MIB (ps.mib) to Enabled (1).
4464	Occasionally, data reported by the frdump and frMgrdump debug-level commands was lost.
4558	During OSPF initialization, link-state database exchange with Wellfleet routers would fail, because the unicast packets transmitted by the MAX TNT had IP TTL set to 0. Could not bring up OSPF between MAX TNT and Wellfleet router.
4595	CLID was not passed across the PRI trunk for net-to-net calls.
4639	When IPDC protocol was used for SS7 signaling between the MAX TNT and a Softswitch and the Softswitch sent RCR message on a channel that was idle, the MAX TNT responded with MRJ (message reject) instead of the expected ACR.
4668	Frame Relay profile values were not case-sensitive.
4670	Outbound modem sessions were not sending the username in the RADIUS Accounting Start record.
4708	Invalid names appeared when xDSL profiles were deleted and recreated.
4770	When using TFTP to load or save a large number of profiles, the TFTP operation could sometimes time out and abort.
4792	After upgrading IPDC server software, the Control Server repeatedly sent a specific sequence of IPDC messages which caused DDL to corrupt its layer 2 transmission window.
4858	When the MAX TNT was set to MAX-compatible Syslog mode, the Incoming Call, Call Connected, and Call Disconnected Syslog records referenced the incoming line's shelf, slot and port values rather than the modem or controller shelf, slot and port as it should for MAX compatibility.
4896	PPP call failed on MAX TNT units, generating Warning 179 messages in Diagnostic mode, when configured with Auth-Frm-Addr-Start set to yes.

TR	Problem corrected
4922	The MAX TNT EchoSuppressToneDisable parameter disabled the tone for all calls.
4946	Analog calls connected to terminal server with digital NAS Port Type.
4966	The MAX TNT did not report the RADIUS Accounting NAS-Port attribute as documented.
4967	Variables were missing for L2TP/ATMP/PPTP in the Ascend MIB.
5113	TDM path was not established correctly for Net-to-Net calls.
5221	The MAX TNT unit incorrectly detected CLID authentication failures because it sent Calling-Station-Id with the incorrect size in RADIUS Access-Request packets.
5308	The MAX TNT unit terminated a call during logon sequence after receiving a Control-C command.
5482	After a few weeks, synchronous users could not connect successfully into the MAX TNT unit. Diagnostic mode displayed the message CANT OPEN 1 HDLC CHANNEL.
5539	MAX TNT unit sent RADIUS authentication and accounting packets from source port 0 to destination port 0.
5626	The SNMP object ssnActiveIdleTime was not supported.
5674	Occasionally, when a Lucent Winmodem connected to a CSM3V modem, it failed to train at speeds lower than V.90, causing the modems to continue to train until timeout.
1000033	MAX TNT units sent RADIUS authentication and accounting packets from source port 0 to destination port 0.
1000235	MAX TNT units supporting E1/R2 did not send blocking towards switch when DSP resources were unavailable.
1000236	MAX TNT units did not support R2 blocking for all R2 types.
1000268	The Series56 III slot card reset, generating a Warning 106 in Diagnostic mode.
250094	Some RADIUS accounting records showed zero packet and byte counts.
250283	Using different trunk numbers for each BRI channel does not work.
258625	The MAX TNT unit did not send SNMP Accounting Event messages correctly.
258669	If you set the Shared-Profile parameter in the IP-Global profile to No, the system always rejected an attempt to share a user profile, even if the RADIUS user profile specified Ascend-Shared-Profile-Enable=Yes.
258671	If you set the Shared-Profile parameter in the IP-Global profile to No, the system always rejected an attempt to share a user profile, even if the RADIUS user profile specified Ascend-Shared-Profile-Enable=Yes.
258749	Inverse ARP was not recognized as a valid protocol on the DS3-ATM card.

Corrections in MAX TNT TAOS 8.0.0

TR	Problem corrected
258791	When the MAX TNT was configured as both ATMP Home Agent and Foreign Agent, and the home network was unreachable, the Foreign Agent did not build a tunnel to the secondary Home Agent.
258800	If an incoming call was received while a modem was dropping a previous call, some of the modems occasionally stopped answering calls.
258821	After an several days, ATMP Gateway Home Agents stopped accepting tunnel requests from Foreign Agents.
258826	Although several tunnels might be active, the Home Agent (in router mode) displayed only the latest tunnel when browsing the ATMP MIB Tunnel Table.
258864	With the MAX TNT configured for Net5 ISDN signalling (framing 2DS) and configured for SNMP, when the MAX TNT unit's D channel failed and recovered from the failure (the line goes to TE) the unit did not send a link up TRAP.
258929	When supporting MultiVoice, an ingress MAX TNT unit sent the ALERTING message before the egress unit received it.
1000079	When an E1 line was incorrectly configured to support Australia PRI (when a Net5 PRI should have been configured), the MAX TNT unit reset, generating a Fatal Error 8 in Diagnostic mode.
1000103	When a client used the Finger command to query a MAX TNT, the connection was refused.
1000135	Blocking needed for NZ IHUG when channels are disabled.
1000192	Systems configured for E1 R2 signaling in China did not pass the correct CLID number or remove the group-II tone.
1000235	MAX TNT did not send blocking messages to the switch when DSP resources became unavailable.
1000236	MAX TNT did not send blocking messages to the switch when DSP resources became unavailable.