MAX Glossary

Ascend Communications, Inc.
Part Number: 7820-0648-001
For Software Version 7.0.0

Ascend Customer Service

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

Obtaining technical assistance

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

Enabling Ascend to assist you

If you need to contact Ascend 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 Ascend product.
- Type of computer you are using.
- Description of the problem.

Calling Ascend from within the United States

In the U.S., you can take advantage of Priority Technical Assistance or an Ascend Advantage Pak 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-ASND (2763) to reach Ascend's 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.

Ascend Advantage Pak

Ascend Advantage Pak is a one-year service contract that includes overnight advance replacement of failed products, technical support, software maintenance releases, and software update releases. For more information, call (800) ASCEND-4 (272-3634), or access Ascend's Web site at `www.ascend.com` and select Services and Support, then Advantage Service Family.

Other telephone numbers

For a menu of Ascend's services, call (800) ASCEND-4 (272-3634). Or call (510) 769-6001 for an operator.

Calling Ascend from outside the United States

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

Telephone outside the United States	(510) 769-8027
Asia Pacific (except Japan)	(+61) 3 9656 7000
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

Note: For additional support information for the Asia Pacific Region, refer to `http://apac.ascend.com/contacts.html`.

Obtaining assistance through correspondence

Ascend 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 Ascend's U.S. offices. Following are the ways in which you can reach Ascend 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 Ascend at the following address:

  Attn: Customer Service
  Ascend Communications, Inc.
  One Ascend Plaza
  1701 Harbor Bay Parkway
  Alameda, CA 94502-3002

Finding information and software on the Internet

Visit Ascend's Web site at `http://www.ascend.com` for technical information, product information, and descriptions of available services.

Visit Ascend's FTP site at `ftp.ascend.com` for software upgrades, release notes, and addenda to this manual.

About This Guide

How to use this guide

This guide contains definitions of technical terms and acronyms commonly found in Ascend documentation. Use this guide as a reference when installing, configuring, or maintaining your system.

Note: This guide describes the full set of features for the MAX running software version 7.0.0. Some features might not be available with older versions or specialty loads of the software.

Documentation conventions

Ascend uses standard documentation conventions. The introductory section of each manual includes a section that describes the conventions, which are as follows:

Convention	Meaning
`Monospace text`	Represents text that appears on your computer's screen, or that could appear on your computer's screen.
`Boldface mono-space text`	Represents characters that you enter exactly as shown (unless the characters are also in ***`italics`***—see *Italics*, below). If you could enter the characters but are not specifically instructed to, they do not appear in boldface.
Italics	Represent variable information. Do not enter the words themselves in the command. Enter the information they represent. In ordinary text, italics are used for titles of publications, for some terms that would otherwise be in quotation marks, and to show emphasis.
[ ]	Square brackets indicate an optional argument you might add to a command. To include such an argument, type only the information inside the brackets. Do not type the brackets unless they appear in bold type.
|	Separates command choices that are mutually exclusive.
>	Points to the next level in the path to a parameter or menu item. The item that follows the angle bracket is one of the options that appears when you select the item that precedes the angle bracket.
Key1-Key2	Represents a combination keystroke. To enter a combination keystroke, press the first key and hold it down while you press one or more other keys. Release all the keys at the same time. (For example, Ctrl-H means hold down the Control key and press the H key.)
Note:	Introduces important additional information.

Related publications

If you need more background information than the documentation set provides, many external references are readily available on the Web or in technical bookstores. You'll find a partial list of such references below

Related RFCs

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

Information about PPP connections

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

- RFC 1662: *PPP in HDLC-like Framing*
- RFC 1661: *The Point-to-Point Protocol (PPP)*
- RFC 1994: *PPP Challenge Handshake Authentication Protocol (CHAP)*
- RFC 1934: *Ascend's Multilink Protocol Plus (MP+)*
- RFC 1969: *The PPP DES Encryption Protocol (DESE)*
- RFC 1989: *PPP Link Quality Monitoring*
- RFC 1990: *The PPP Multilink Protocol (MP)*
- RFC 2125: *The PPP Bandwidth Allocation Control Protocol (BACP)*
- RFC 2153: *PPP Vendor Extensions*
- RFC 1962: *The PPP Compression Control Protocol (CCP)*
- RFC 1974: *PPP Stac LZS Compression Protocol*
- RFC 1877: *PPP Internet Protocol Control Protocol Extensions for Name Server Addresses*
- RFC 1618: *PPP over ISDN*
- RFC 1332: *The PPP Internet Protocol Control Protocol (IPCP)*
- RFC 1552:T*he PPP Internetwork Packet Exchange Control Protocol (IPXCP)*
- RFC 1378: *The PPP AppleTalk Control Protocol (ATCP)*

Information about IPX routing

For information about IPX routing, see RFC 1634: *Novell IPX Over Various WAN Media (IPXWAN).*

Information about IP routers

RFCs that describe the operation of IP routers include:

- RFC 1812: *Requirements for IP Version 4 Routers*
- RFC 1519: *Classless Inter-Domain Routing (CIDR): An Address Assignment and Aggregation Strategy*
- RFC 2002: *IP Mobility Support*
- RFC 2030: *Simple Network Time Protocol (SNTP) Version 4 for IPv4, IPv6 and OSI*
- RFC 1256: *ICMP Router Discovery Messages*
- RFC 1393: *Traceroute Using an IP Option*
- RFC 1433: *Directed ARP*
- RFC 1582: *Extensions to RIP to Support Demand Circuits*
- RFC 1787: *Routing in a Multi-provider Internet*

Information about OSPF routing

For information about OSPF routing, see:

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

Information about multicast

For information about multicast, see:

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

Information about firewalls and packet filtering

RFCs that describe firewalls and packet filters include:

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

Information about general network security

RFCs pertinent to network security include:

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

Information about external authentication

For information about RADIUS and TACACS authentication, see:

- RFC 2138: *Remote Authentication Dial In User Service (RADIUS)*
- RFC 1492: *An Access Control Protocol, Sometimes Called TACACS*

ITU-T recommendations

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

Related books

The following books are available in technical bookstores.

- *Routing in the Internet*, by Christian Huitema. Prentice Hall PTR, 1995. Recommended for information about IP, OSPF, CIDR, IP multicast, and mobile IP.
- *SNMP, SNMPV2 and RMON: Practical Network Management*, by William Stallings. Addison-Wesley, 1996. Recommended for network management information.
- *Enterprise Networking: Fractional T1 to Sonet Frame Relay to Bisdn*, by Daniel Minoli. Artech House, 1993. Recommended as a WAN reference.
- *TCP/IP Illustrated*, volumes 1&2, by W. Richard Stevens. Addison-Wesley, 1994.

Alphabetic list of terms

Numeric

2DS—A variant of the standard G.703 framing required by most E1 DPNSS providers in the U.K See also *G.703*.

3.1 Khz audio-bearer service—A service that sends a data call over a voice trunk. Because echo cancellation corrupts data transmitted on voice trunks, each switch should turn off echo cancellation on the trunks handling 3.1 Khz audio-bearer service. The 3.1 Khz audio-bearer service is sometimes referred to as *Data Over Subscriber Bearer Service (DOSBS).*

7-bit mode—See *ASCII mode*.

8-bit Binary mode—See *Binary mode*.

10Base-T—The 802.3 IEEE standard for operating a 10-Mbps Ethernet network with twisted-pair cabling and a wiring hub. 10Base-T is also known as *UTP Ethernet* and *twisted-pair Ethernet*. See also *10Base-T hub*.

10Base-T hub—A hub providing a common termination point for hosts connected to 10Base-T wiring. See also *10Base-T*.

10Base-T interface—A MAX interface that supports a 10Base-T connection to a LAN. The 10Base-T interface features autosensing and full-duplex capability. See also *10Base-T*.

56K modem—An analog modem that offers a high-speed alternative to 33.6K modems. However, the increased data rate over the local loop is only possible given the proper network conditions. A true 56K implementation works only on the downstream path of the call, requires an end-to-end 56-Kbps solution, depends on the existence of a single analog-to-digital (A-D) conversion, works over the local loop using the existing infrastructure, and requires a trunk-side connection to the Central Office (CO) switch. See also *A-D conversion*, *downstream path*, *local loop*, *trunk-side connection*.

100Base-T—The 802.3 IEEE standard for operating a 100-Mbps Ethernet network. It differs from the 10Base-T standard by requiring higher-grade cable or more wiring pairs, and by supporting cable lengths that are only a tenth as long as 10Base-T cable lengths. See also *10Base-T*.

100Base-T interface—A MAX interface that supports a 100Base-T connection to a LAN. The 100Base-T interface features autosensing and full-duplex capability. See also *100Base-T*.

802.2—An IEEE protocol specification for the Media Access Control (MAC) header of an IPX frame in NetWare 3.12 or later. An 802.2 frame contains the Logical Link Control (LLC) header in addition to the MAC header. Compare with *802.3*, *Ethernet II*, *SNAP*. See also *IPX frame, LLC, MAC*.

802.3—An IEEE protocol specification for the Media Access Control (MAC) header of an IPX frame in NetWare 3.11 or earlier. An 802.3 frame does not contain the Logical Link Control (LLC) header in addition to the MAC header. The 802.3 frame is also called *Raw 802.3*. Compare with *802.2*, *Ethernet II*, *SNAP*. See also *IPX frame, LLC, MAC*.

802.5—An IEEE protocol specification for the physical layer and Media Access Control (MAC) sublayer of a token-ring topology. 802.5 implements token passing over Shielded Twisted Pair (STP) cabling, and offers data rates of 4 or 16 Mbps. See also *STP cable*.

A

ABR—Area Border Router. An ABR is an Open Shortest Path First (OSPF) router that belongs to both a regular area and the backbone area. See also *area*, *backbone area*, *OSPF*.

Access-Accept packet—A packet sent by the RADIUS server to inform the MAX that a client's request for access has been granted. See also *RADIUS server*.

Access-Challenge packet—A request for the user to enter a password in a hand-held token card. The token-card server sends the Access-Challenge packet through the RADIUS server and the MAX to the user. See also *RADIUS server*, *token card*, *token-card server*.

access concentrator—A device that efficiently forwards data, handling incoming calls for a network Point Of Presence (POP). In general, an access concentrator supports dial-in modem calls, ISDN connections, nailed-up links, Frame Relay traffic, and multiprotocol routing. See also *dial-in modem access*, *Frame Relay Direct*, *ISDN*, *nailed-up circuit*, *POP*.

Access Control—See *Ascend Access Control*.

Access-Password-Ack packet—A response from the RADIUS server informing the MAX that it has accepted a new password. See also *RADIUS server*.

Access-Password-Reject packet—A response from the RADIUS server informing the MAX that it has rejected a new password. See also *RADIUS server*.

Access-Password-Request packet—A password-change request that the MAX sends to the RADIUS server. See also *RADIUS server*.

Access-Reject packet—A packet the RADIUS server sends to inform the MAX that it has not granted a client's request for access. The RADIUS server sends an Access-Reject packet if the user enters an unknown user name, fails to enter the correct password, or enters an expired password. See also *RADIUS server*.

Access-Request packet—A packet that the MAX sends to the RADIUS server on behalf of a client attempting to establish a connection. See also *RADIUS server*.

accounting—A way to log information in RADIUS about Start session, Stop session, and Failure-to-start session events. When the MAX recognizes one of these events, it sends an accounting request to RADIUS. When the accounting server receives the request, it combines the information into a record and timestamps it. Each type of accounting record contains attributes associated with an event type, and can show the number of packets the MAX transmitted and received, the protocol in use, the user name and IP address of the client, and other session information. See also *accounting server*, *Failure-to-start session*, *Start session*, *Stop session*.

Accounting-Request—A request for accounting information. The MAX sends an Accounting-Request packet to the RADIUS accounting server. See also *accounting server*, *RADIUS*.

Accounting-Response—A packet containing accounting information. The RADIUS accounting server sends an Accounting-Response packet to the MAX. See also *accounting server*, *RADIUS*.

accounting server—The RADIUS daemon with accounting enabled. You can specify that all users access the same accounting server. Or, you can specify access on a per-user basis. See also *accounting*, *RADIUS daemon*.

ACE authentication—A form of token-card authentication in which RADIUS forwards a connection request to a Security Dynamics ACE/Server. The ACE/Server sends an Access-Challenge packet back through the RADIUS server and the MAX to the user dialing in. The user sees the challenge message, obtains the current token from his or her card, and enters the token. (A token is a type of password.)

The token travels back through the MAX and the RADIUS server to the ACE/Server. The ACE/Server sends a response to the RADIUS server, specifying whether the user has entered the proper user name and token. If the user enters an incorrect token, the ACE/Server returns another challenge, and the user can again attempt to enter the correct token. The server sends up to three challenges. After three incorrect tries, the MAX terminates the call.

ACE authentication is also known as *SecurID authentication*.

See also *ACE token*, *authentication*, *RADIUS server*, *token*, *token card*, *token-card authentication*, *token-card server*.

ACE token—A randomly generated access code that a user obtains from a SecurID token card. The code changes every 60 seconds. See also *ACE authentication*, *token card*.

ACK—Acknowledgment. An ACK is a packet that the system uses to acknowledge a transmission. When a device receives a packet, it sends back a packet to the sending device. If all the data arrived without corruption, the packet is an ACK. If some of the data is missing or corrupt, a NAK results, and acts as a request that the sender retransmit the data.

acknowledgment—See *ACK*.

A-D conversion—Analog-to-digital conversion. A-D conversion is a process in which an analog signal is modified into a digital signal. A-D conversion takes place, for example, when an analog modem call reaches a digital modem. Compare with *D-A conversion*. See also *analog signal*, *digital modem*, *digital signal*.

add-on number—One or more numbers that the MAX uses to build a multichannel MP, MP+, AIM, or BONDING call. A multichannel call begins as a single-channel connection to one phone number. The calling unit then requests additional phone numbers to connect additional channels, and stores the add-on numbers it receives from the answering unit. To add channels to the call, the calling unit must integrate the add-on numbers with the phone number it dialed initially. Typically, the phone numbers assigned to the channels share a group of leading digits. The add-on numbers are the rightmost digits identifying each phone number, excluding the digit(s) that the phone numbers have in common.

address resolution—A method of mapping a logical address (such as an IP address) to a hardware address (such as a MAC address). See also *ARP*, *hardware address*, *IP address*, *logical address*, *MAC address*.

Address Resolution Protocol—See *ARP*.

adjacency—A relationship formed between neighboring Open Shortest Path First (OSPF) routers for the purpose of exchanging routing information. An OSPF router dynamically detects its neighboring routers by sending Hello packets to the multicast address AllSPFRouters. It then attempts to form adjacencies (Figure 1).

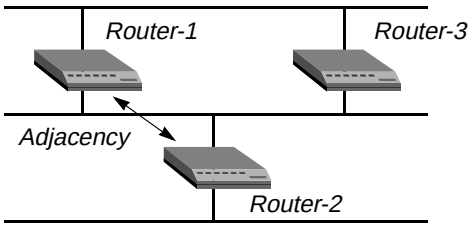

Figure 1. Adjacency between neighboring routers

Neighbors exchange databases and build a consistent, synchronized database between them. When an OSPF router detects a change on one of its interfaces, it modifies its link-state database and multicasts the change to its adjacent neighbor, which in turn propagates the change to its adjacent neighbor, until all routers within an area have synchronized link-state databases. This method of updating routing information results in quick convergence among routers.

See also *area*, *convergence*, *link-state database*, *OSPF*, *router*.

Advanced Mobile Phone System—See *AMPS*.

AEP—AppleTalk Echo Protocol. AEP is a Transport-layer protocol that enables the network to determine whether two nodes are connected and available to receive transmissions.

agent—A network device (such as the MAX) that provides Simple Network Management Protocol (SNMP) information to a manager application running on another computer. The agent and manager share a database of information, called the *Management Information Base (MIB)*. The manager polls the agent for information at regular intervals. When an unusual system event occurs, the agent can use a message called a *traps-PDU* to send unsolicited information to the manager. See also *manager*, *MIB*, *SNMP*, *traps-PDU*.

AIM—Ascend Inverse Multiplexing. Developed and supported by Ascend Communications, AIM manages the connection of two remotely located inverse multiplexers. See also *inverse multiplexer*, *inverse multiplexing*.

AIM/6 card—See *Host/6 card*.

AIM port—A MAX port that supports AIM and BONDING on a Dual/Host or Host/6 card. See also *dual IP*, *Host/6 card*.

AIS—Alarm Indication Signal. An AIS is a signal that a device sends in order to take a T1 line out of service. See also *T1 line*.

alarm—A signal that indicates that the system has detected a security violation or error. See also *alarm event*.

alarm event—A type of alarm. The following alarm events are defined in the Ascend Enterprise MIB:

Alarm event	Indicates that the MAX unit
coldStart (RFC-1215 trap-type 0)	Is reinitializing itself in such a way that it might alter the configuration of either the SNMP manager or the unit.
warmStart (RFC-1215 trap-type 1)	Is reinitializing itself so that neither the configuration of the SNMP manager nor that of the unit will change.
linkDown (RFC-1215 trap-type 2)	Recognizes a failure in one of the communication links represented in the SNMP manager's configuration.
linkUp (RFC-1215 trap-type 3)	Recognizes that one of the communication links represented in the SNMP manager's configuration has come up.
frDLCIStatusChange (RFC-1315 trap-type 1)	Recognizes that one of the Virtual Circuits (VCs) has changed states. The link has been created, invalidated, or toggled between the active and inactive states.
eventTableOverwrite (Ascend trap-type 16)	Detected that a new event has overwritten an unread event. The unit sends this trap only for systems that support Ascend's accounting MIB. Once sent, additional overwrites will not cause another trap to be sent until at least one table's worth of new events has occurred.

See the Ascend Enterprise MIB for the most up-to-date information.

Alarm Indication Signal—See *AIS*.

alarm relay—A mechanism whose contacts remain open on the MAX unit's back panel's alarm relay terminal block during normal operation. If you enable them, the alarm relay contacts close during loss of power, hardware failure, or a system reset. You can also specify whether the contacts close when the bit-error rate exceeds a certain value, or when all T1 PRI lines go out of service.

A-Law—An ITU-T standard for sampling data by means of Pulse Coded Modulation (PCM). A-Law is most commonly used outside of North America and Japan. Compare with *U-Law*. See also *PCM*.

Alternate Mark Inversion—See *AMI*.

ALU—Average Line Utilization. ALU is the average amount of bandwidth used on a line over a user-specified period of time. The MAX uses ALU when determining whether to add or subtract bandwidth from a multichannel call. See also *DBA*.

Always On/Dynamic ISDN—See *AO/DI*.

American National Standards Institute—See *ANSI*.

AMI—Alternate Mark Inversion. An encoding method in which alternating positive and negative voltage represents a 1, and zero voltage represents a zero.

AMPS—Advanced Mobile Phone Service. AMPS is a standard system for analog cellular telephone service. Introduced by AT&T in 1983, AMPS is the most widely used cellular system in the United States. The service uses frequency ranges between 800 and 900 MHz. Each provider can use half of the 824-849 MHz range for receiving signals, and half the 869-894 MHz range for transmission.

analog data—Data that can change continuously and have any value in a range. Examples of analog data are the time of day represented by clock hands, and the temperature represented by a liquid thermometer. Compare with *digital data*. See also *analog signal*.

analog line—A line that transmits data by means of an analog signal. Compare with *digital line*. See also *analog signal.*

analog loopback—A test that checks whether the modem or Data Terminal Equipment (DTE) is causing errors in data transmission. During an analog loopback, the system sends data between the local modem and the local DTE. Errors in transmission indicate a problem with the modem, DTE, or the interface between them. Compare with *digital loopback*. See also *local loopback, loopback, remote loopback*.

analog signal—A type of signal that encodes data transmitted over wire or through the air, commonly represented as an oscillating wave. An analog signal can transmit analog or digital data. It takes any value in a range, and changes smoothly between values. A radio station sends analog music data using analog signals, while a modem transmits digital data using analog signals. Compare with *digital signal*. See also *analog data*.

analog-to-digital conversion—See *A-D conversion*.

ANI—Automatic Number Identification. ANI is a mechanism that informs the called party of the calling party's phone number. Though ANI is often thought of as an ISDN feature, it is actually part of Signaling System 7, and distinct from ISDN. See also *CLID authentication, Signaling System 7*.

Annex A—See *Frame Relay Annex A.*

Annex D—See *Frame Relay Annex D*.

ANSI—American National Standards Institute. ANSI creates standards for networking and communications. It is the U.S. representative to the International Standards Organization (ISO). See also *ISO*.

ANSI T1.617 Annex D—See *Frame Relay Annex D.*

answer number—A phone number used for call-routing purposes. It appears in a number of profiles. In each case, it indicates "route calls received on this number to me." For example, an answer number specified in the Ethernet profile indicates that calls received on that number should be routed to the bridge/router. In a Modem profile, the answer number indicates that calls received on that number should be routed to an available digital modem. See also *bridge, call routing, digital modem, router*.

Answer profile—A profile that sets baseline values to determine how the MAX evaluates incoming calls before it accepts them. If the call does not comply with the Answer profile settings, the unit rejects the call without answering it. Therefore, you must check the Answer profile values to make sure they are appropriate for your site.

The MAX applies the Answer profile values before it routes the call and locates a Connection profile or RADIUS user profile. If the caller's profile contains a parameter or attribute similar to one in the Answer profile, but the caller's setting specifies a different value, the MAX uses the value in the Connection profile or RADIUS user profile to build the session.

By default, the Answer profile enables all types of encapsulation and routing, and the basic call-setup parameters use the lowest-common-denominator settings. The default settings are appropriate for many sites. You might want to change the settings in order to finetune the criteria by which the MAX accepts calls or determines how much bandwidth is accessible to Multilink Protocol (MP) or Multilink Protocol Plus (MP+) sessions. See also *Connection profile*, *MP*, *MP+*.

AO/DI—Always On/Dynamic ISDN. AO/DI is a networking service that enables you to send and receive data through a nailed-up X.25 connection (supported over an ISDN D channel, ISDN B channel, or nailed-up 56K line). The MAX uses switched ISDN B channels only when required, on the basis of increased bandwidth use. Through its use of AO/DI, X.25, and Bandwidth Allocation Control Protocol (BACP), the MAX avoids dialup charges and the use of switched B channels whenever it sends or receives data over the X.25 connection.

In a traditional ISDN environment, data moves across B channels, and signaling information moves across the D channel. Because signaling information uses a small percentage of available D-channel bandwidth, AO/DI was developed to maximize bandwidth usage while reducing the necessity that all data travel over B channels.

Among the functions that can use AO/DI are the following:

- Transfer of email
- Reception of news broadcasts
- Automated collection of data

For all Ascend units, AO/DI enables you to use X.25 bandwidth of up to 9600 bps. If data transfers require more bandwidth, the MAX dials and combines B channels by means of BACP. See also *BACP*, *B channel*, *D channel*, *X.25*.

APP—Ascend Password Protocol. APP is a User Datagram Protocol (UDP) that enables a user to respond to password challenges received from an external authentication server. See also *AppleTalk*.

AppleTalk—Apple's protocol suite that enables Macintosh computers to function on a network. AppleTalk works with such network operating systems as TOPS (from Sun Microsystems) and AppleShare. See also *AppleTalk router*, *AppleTalk routing*, *ARA*.

AppleTalk Control Protocol—See *ATCP*.

AppleTalk Echo Protocol—See *AEP*.

AppleTalk Remote Access—See *ARA*.

AppleTalk Remote Access Protocol—See *ARAP*.

AppleTalk router—A device that sends AppleTalk packets from a source to a destination by various paths. See also *AppleTalk*, *AppleTalk routing*, *ARA*.

AppleTalk routing—A routing configuration in which Macintosh computers can share files and services on a network. A MAX configured for AppleTalk routing can receive dial-in connections from AppleTalk Remote Access (ARA) client software, Point-to-Point Protocol (PPP) dial-in software that supports AppleTalk, and AppleTalk-enabled Ascend units.

Figure 2 shows a MAX that routes AppleTalk between WAN interfaces and a local AppleTalk interface.

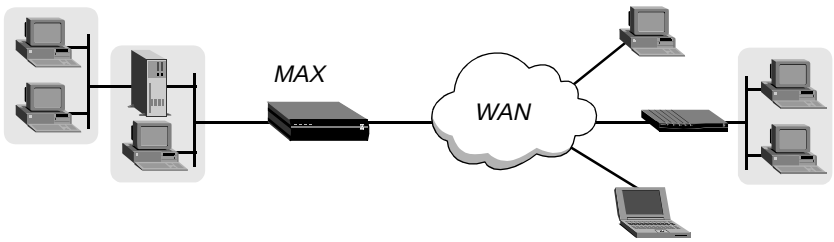

Figure 2. Routing AppleTalk between LAN and WAN interfaces

PPP and ARA are the encapsulation protocols used for AppleTalk dial-in connections. ARA 3.0 supports both ARA and PPP. You can use AppleTalk PPP and ARA over a modem or V.120 ISDN TA connection. You can also use AppleTalk PPP over synchronous PPP when the calling unit is an Ascend router.

You must enable AppleTalk routing for any kind of AppleTalk connection, even if the individual connection to a remote device does not use routing. See also *AppleTalk*, *AppleTalk router*, *ARA*, *PPP*.

Application layer—The highest layer of the OSI Reference Model. The Application layer provides applications with access to the network. File transfer, email, and network management software are examples of Application-layer programs. Protocols such as File Transfer Protocol (FTP), Rlogin, Simple Network Management Protocol (SNMP), and Telnet provide Application-layer services. See also *FTP*, *OSI Reference Model*, *Rlogin*, *SNMP*, *Telnet*.

APP Server utility—Ascend Password Protocol Server utility. The APP Server utility enables a user to respond to password challenges received from an external authentication server, such as an ACE/Server or SafeWord server. To allow a user to supply a password from a host on the local network, you must configure the MAX to communicate with the APP Server utility on that host. See also *authentication server*, *token-card server*.

ARA—AppleTalk Remote Access. ARA enables a remote Macintosh workstation to gain access to an IP network. You can use ARA over a modem or V.120 connection, or over synchronous PPP when the calling unit is an AppleTalk-enabled Ascend unit. Clients can dial in using ARA client software or a PPP dialer that supports AppleTalk. See also *AppleTalk*, *AppleTalk router*, *AppleTalk routing*, *modem*, *PPP*, *V.120*.

ARAP—AppleTalk Remote Access Protocol. ARAP is an AppleTalk protocol that enables a remote Macintosh or PowerBook computer to connect to a LAN. See also *AppleTalk*.

ARCnet—Attached Resource Computer Network. ARCnet is a baseband network architecture with a transmission rate of up to 2.5 Mbps. Because it is relatively inexpensive and easy to set up, ARCnet is typically used for smaller networks.

area—A portion of an Open Shortest Path First (OSPF) Autonomous System (AS). An area acts as its own network. All area-specific routing information stays within the area, all routers within an area have a synchronized link-state database, and each database within an area is unique. On the MAX, an area number uses dotted decimal notation. It is not an IP address.

To tie the areas together, some routers belong to a backbone area and one other type of area. These routers are called *Area Border Routers (ABRs)*. In Figure 3, all of the routers are ABRs.

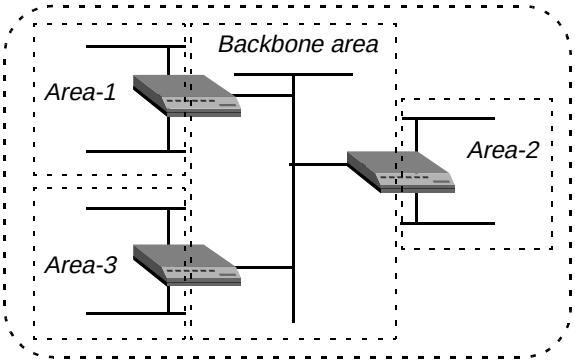

Figure 3. Dividing an AS into areas

See also *ABR*, *AS*, *backbone area*, *link-state database*, *normal area*, *NSSA*, *OSPF*, *router*, *stub area*.

Area Border Router—See *ABR*.

ARP—Address Resolution Protocol. ARP is a protocol in the TCP/IP protocol suite. By mapping an IP address to a physical (hardware) address, ARP enables a unit to identify hosts on an Ethernet LAN. See also *Ethernet*, *proxy ARP*, *TCP/IP*.

AS—Autonomous System. An AS is a group of Open Shortest Path First (OSPF) routers that exchange information, typically under the control of one company. An AS can include a large number of networks, all of which share the same AS number. All information exchanged within the AS is interior. Exterior protocols, such as Exterior Gateway Protocol (EGP), exchange routing information between one AS and another. Using an EGP, the MAX imports external routes into its OSPF database and flags them as ASE (Autonomous System External). See also *ASE*, *EGP*, *external route*, *OSPF*, *router*.

ASBR—Autonomous System Border Router. An ASBR is an Open Shortest Path First (OSPF) router that handles communication between Autonomous Systems (AS) by using an Exterior Gateway Protocol (EGP), as shown in Figure 4.

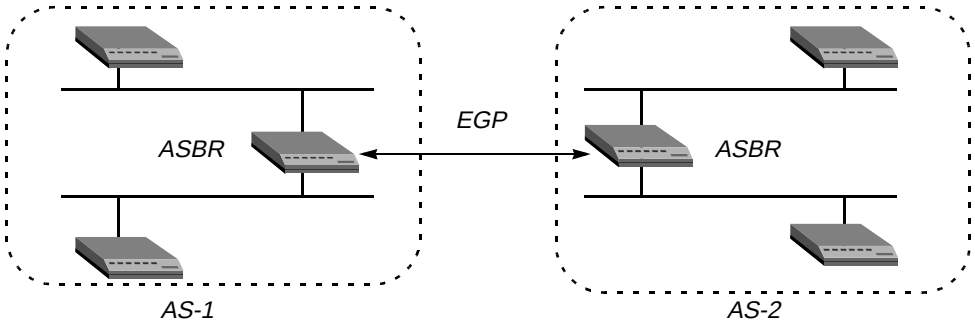

Figure 4. Autonomous System Border Routers (ASBRs)

ASBRs perform calculations related to external routes. The MAX imports external routes from Routing Information Protocol (RIP)—for example, when it establishes a WAN link with a caller that does not support OSPF—and the ASBR calculations are always performed.

Compare with *ABR*. See also *AS*, *EGP*, *external route*, *OSPF*.

Ascend Access Control— A comprehensive network-wide security-management system based on industry-standard RADIUS. Ascend Access Control enables you to identify legitimate callers, perform authentication and authorization, monitor access to network resources, and compile extensive billing and accounting details. See also *accounting*, *authentication*, *authorization*, *RADIUS*.

Ascend-Access-Event-Request packet— A packet containing a notification that the MAX has started up, or making a request for the RADIUS server to record the number of open sessions. See also *RADIUS server*.

Ascend-Access-Event-Response packet—A response from the RADIUS server reporting that the MAX has started up, or specifying the number of open sessions and informing the MAX that the server has received and recorded the MAX unit's ID. See also *RADIUS server*.

Ascend-Access-New-Pin packet—A response from the RADIUS server informing the MAX that it should request access again, but with the next Personal Identification Number (PIN) in the sequence. See also *RADIUS server*.

Ascend-Access-Next-Code packet—A response from the RADIUS server informing the MAX that it should request access again, but with the next password in the sequence. See also *RADIUS server*.

Ascend Enterprise Traps MIB—A MIB containing specifications for traps that denote important events. The trap classifications are:

- Error events

  The error events reported are those defined in RFC 1215. They are `coldStart`, `warmStart`, `linkDown`, and `linkUp`.

- Port-state-change events

  Port-state-change events report changes in activity at a port, such as when a call goes from active to inactive.

- Security events

  Security events are used to notify users of security problems and track access to the unit from the console.

See also *MIB*.

Ascend Inverse Multiplexing—See *AIM*.

Ascend-Password-Expired packet—A response from RADIUS server to the MAX, indicating that the password the user entered matches the one in the user profile, but has expired. (That is, the Access-Request packet sent a valid but expired password.) See also *RADIUS server*.

Ascend-Terminate-Session packet—A response from the RADIUS server informing the MAX that it should terminate the session and display the message sent in the packet. See also *RADIUS server*.

Ascend Tunnel Management Protocol—See *ATMP*.

ASCII—American Standard Code for Information Interchange. ASCII is a character-encoding system for Local Area Networks (LANs). The 128 standard ASCII characters are composed of seven bits, and have the values 0–127. The extended ASCII character set contains another 128 values.

ASCII users file—See *flat ASCII users file*.

ASCII mode—A Telnet mode for terminal-server users. In ASCII mode, bit 8 is set to 0 (zero). ASCII mode is also called *standard 7-bit mode* or *Network Virtual Terminal (NVT) ASCII*. This mode is the default if no other mode is specified. Compare with *Binary mode*, *Transparent mode*. See also *Telnet*, *Telnet mode*.

ASE—Autonomous System External. The MAX uses the term ASE to denote external routes it imports into its Open Shortest Path First (OSPF) database. The MAX redistributes these routes by means of OSPF ASE advertisements, and propagates its OSPF routes to remote WAN routers running Routing Information Protocol (RIP). See also *external route*, *OSPF*, *RIP*, *router*.

ASE Type-5—Autonomous System External Type-5. ASE Type-5 is an external route originated by an Area Border Router (ABR) as a Link-State Advertisement (LSA). An Open Shortest Path First (OSPF) normal area allows Type-5 LSAs to be flooded throughout it.

A Not So Stubby Area (NSSA) and a stub area do not receive or originate Type-5 LSAs. However, for NSSAs, all routes imported to OSPF have the P-bit set (P stands for *propagate*). When the P-bit is enabled, ABRs translate Type-7 LSAs to Type-5 LSAs, which can then be flooded to the backbone. These external routes are considered Type-7 LSAs.

Compare with *ASE Type-7*. See also *ABR*, *AS*, *ASE*, *external route*, *LSA*, *normal area*, *NSSA*, *OSPF*, *stub area*.

ASE Type-7—Autonomous System External Type-7. ASE Type-7 is a type of Link-State Advertisement (LSA) defined for Not So Stubby Areas (NSSAs) in Open Shortest Path First (OSPF) version 2. For NSSAs, all routes imported to OSPF have the P-bit set (P stands for *propagate*). When the P-bit is enabled, ABRs translate Type-7 LSAs to Type-5 LSAs, which can then be flooded to the backbone. These external routes are considered Type-7 LSAs. Compare with *ASE Type-5*. See also *AS*, *ASE*, *LSA*, *NSSA*, *OSPF*, *stub area*.

ASN.1—Abstract Syntax Notation One. In the OSI Reference Model, ASN.1 is a notation for describing data structures on a network. It provides a consistent syntax for transferring data between different systems. See also *OSI Reference Model*.

asynchronous PPP—A mode for sending Point-to-Point Protocol (PPP) packets. In asynchronous mode, the characters that form the data packets are sent at irregular intervals, without a clocking signal to time transmission. Figure 5 illustrates a single-channel asynchronous PPP call in which the calling device is a modem.

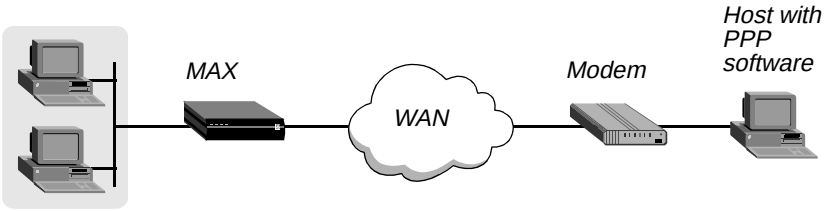

Figure 5. Asynchronous PPP connection

Asynchronous PPP is commonly used in lower-speed transmission and less-expensive transmission systems. Asynchronous PPP is also known as *async PPP.* Compare with *synchronous PPP*. See also *asynchronous transmission*, *PPP*.

Asynchronous Transfer Mode—See *ATM*.

asynchronous transmission—A mode in which the sending and receiving serial hosts know where a character begins and ends because each byte is framed with additional bits, called a start bit and a stop bit. A start bit indicates the beginning of a new character. It is always 0 (zero). A stop bit marks the end of the character. It appears after the parity bit, if one is in use.

An asynchronous link uses the type of serial communication provided by a PC COM port. A dial-up modem or V.120 Terminal Adapter (TA) initiates an asynchronous host-to-network or host-to-host connection. The call can use Point-to-Point Protocol (PPP) encapsulation, V.120 encapsulation, or raw (unencapsulated) Transport Control Protocol (TCP).

The MAX routes an asynchronous call to a digital modem as a voice call, and then to the terminal-server software. If the terminal server does not detect a PPP packet, it begins a login sequence. If the terminal server detects a PPP packet, it passes the call on to the router, where it is handled as a regular PPP connection. The caller never sees the terminal-server interface.

See also *asynchronous PPP*, *digital modem*, *PPP*, *TCP*, *terminal server*, *V.120*, *V.120 TA*.

async PPP—See *asynchronous PPP*.

AT command set—A set of commands created by Hayes Microcomputer Products for operation of its modems. AT stands for the *Attention* signal that precedes each modem command. For example, the command string ATDT enables a modem to dial a number on a Touch Tone system. Most modem manufacturers use the AT command set. It is a de facto industry standard. See also *modem*.

ATCP—AppleTalk Control Protocol. A protocol that enables you to route AppleTalk packets that are encapsulated in Point-to-Point Protocol (PPP).

ATM—Asynchronous Transfer Mode. ATM is a packet-switched, broadband network architecture central to Broadband ISDN (B-ISDN). It provides very high bandwidth, enabling data, voice, and multimedia transmissions to occupy the same line. ATM is also known as *cell relay.* See also *B-ISDN*, *broadband*, *packet switching*.

ATMP—Ascend Tunnel Management Protocol. ATMP provides a tunneling mechanism between two Ascend units across the Internet or a Frame Relay network. Each Ascend unit can be a MAX or a Pipeline 400. The protocol uses standard Generic Routing Encapsulation (GRE) and is based on the User Datagram Protocol (UDP) and Internet Protocol (IP).

ATMP provides a Virtual Private Network (VPN) solution over the backbone resources of Internet Service Providers (ISPs) and carriers. Without ATMP, each Mobile Client and remote user has to dial directly into the network, resulting in long-distance charges. With ATMP, users can make a local call and have the transmission securely tunneled.

Figure 6 shows an ATMP tunnel between two MAX units. A Mobile Client, such as a traveling salesperson, initiates the connection. The unit that authenticates the Mobile Client is the ATMP Foreign Agent. The unit that accesses the Home Network is the ATMP Home Agent. The Home Network is the destination network for Mobile Clients. In Figure 6, the Mobile Client is a salesperson who logs into an ISP (the Foreign Agent) to access his or her Home Network.

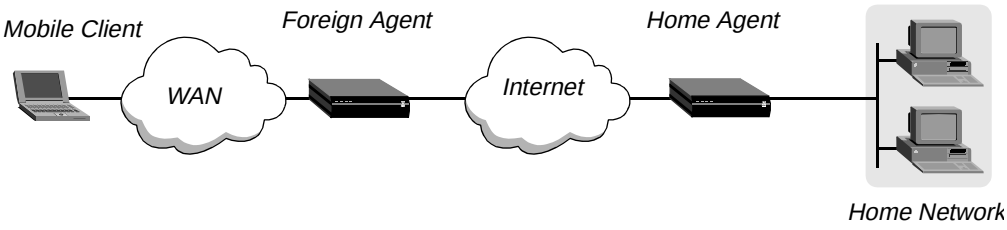

Figure 6. ATMP tunnel across the Internet

As described in RFC 1701, GRE hides packet contents and enables transmission of packets that the Internet would otherwise not accept. When you use ATMP with the MAX, you can transmit either IP packets that use unregistered addresses or IPX packets from roaming clients. See also *Foreign Agent*, *Frame Relay*, *GRE*, *Home Agent*, *Home Network*, *IP*, *IPX*, *ISP*, *Mobile Client*, *UDP*, *VPN*.

Attached Resource Computer Network—See *ARCnet*.

attenuation—The reduction in the strength of a signal over distance, expressed in decibels per kilometer (dB/Km) or per 100 feet. Factors affecting attenuation are the frequency range of the signal, wire shielding, and type of cable. Unshielded Twisted Pair (UTP) cable suffers from the most attenuation, while fiber-optic cable has very little attenuation. See also *attenuator*, *UTP cable*.

attenuator—A device that reduces the amplitude of a signal. See also *attenuation*.

attribute—A series of values in a RADIUS user profile or pseudo-user profile. The attributes indicate a user name and password, and enable you to configure routing, bridging, call management, and usage restrictions. See also *pseudo-user profile*, *RADIUS*, *RADIUS server*, *user profile*.

AUI—Auxiliary Unit Interface. An AUI is a 15-pin D-type connector for Ethernet connections. It typically links a cable to a Network Interface Card (NIC). An AUI is also known as a *Digital, Intel, Xerox (DIX) connector.* See also *Ethernet*, *NIC*.

authentication—A method of identifying users permitted to access network resources. Authentication is the first line of defense against unauthorized access to your network. The MAX supports a variety of authentication methods. You can use:

- Calling-Line ID (CLID) to verify that the call is being placed from a trusted telephone number.
- Called-number authentication to verify the number called.
- Callback security. After authentication is complete, the MAX can hang up and call back, ensuring that the connection is made only with a trusted number.
- Expect-send scripts to authenticate logins to the terminal server.
- Name and password authentication of Point-to-Point Protocol (PPP) calls. The MAX supports Password Authentication Protocol (PAP), PAP with encryption (PAP-DES), Challenge Handshake Authentication Protocol (CHAP), and Microsoft's extension of CHAP (MS-CHAP).
- Token cards. Using a token-card server, you can accept or reject calls by means of PAP-Token, PAP-Token-CHAP, or Cache-Token authentication.

When the MAX is shipped from the factory, it is set to not require any authentication.

See also *Cache-Token authentication*, *called-number authentication*, *CHAP*, *CLID authentication*, *expect-send script*, *PAP*, *PAP-Token authentication*, *PAP-Token-CHAP authentication*, *token card*, *token-card authentication*, *token-card server*.

authentication key—A shared secret passed between the MAX and an authentication server.

- If the MAX is acting as a RADIUS, TACACS, or TACACS+ client, the authentication key is a password supplied by the MAX to the server.
- If the MAX is acting as a Defender client, the authentication key is a DES secret key shared between the MAX and the Defender authentication server. This key is also used for authentication by the MAX in its role as a Defender authentication agent.
- In Open Shortest Path First (OSPF) routing, the authentication key is a 64-bit clear password inserted into the OSPF packet header. It is used by OSPF routers to allow or exclude packets from an area.

See also *authentication server*, *OSPF*, *RADIUS*, *TACACS*, *TACACS+*.

authentication request—A request that the MAX sends an authentication server on behalf of a client requesting access. See also *authentication response*, *authentication server*,

authentication response—A response from an authentication server, notifying the MAX that a user's request for access has been either granted or denied. See also *authentication request*, *authentication server*.

authentication server—An external server, such as a RADIUS, TACACS, TACACS+, or token-card server, that verifies whether a user requesting access to the network has permission to use network resources. See also *RADIUS*, *RADIUS server*, *TACACS*, *TACACS+*, *token-card server*.

authentication timeout—The number of seconds between retries to the external authentication server.

If the MAX is acting as a RADIUS, TACACS, or TACACS+ client, the MAX waits the specified number of seconds for a response to an authentication request. If it does not receive a response within that time, it times out and sends the authentication request to the next authentication server.

If the MAX is acting as a Defender or SecurID client (which support only one server address), the MAX waits the number of seconds you specify before assuming that the server has become nonfunctional.

See also *authentication server*, *RADIUS*, *RADIUS server*, *TACACS*, *TACACS+*.

authenticator field—In a RADIUS packet, a field that enables the system to authenticate packets between the MAX and the authentication server. The MAX and the authentication server share a secret that the system uses, along with the authenticator field, to provide password encryption and packet authentication. The shared secret resides in the `clients` file on the authentication host.

The MAX checks all authentication and accounting packets to ensure that they come from known sources. The check makes use of the shared secret, the authenticator field, and MD5 encoding. In addition, all passwords that the MAX sends are encrypted with MD5, CHAP, or DES. Passwords that the authentication server sends can be encrypted with MD5. See also *authentication*, *CHAP*, *DES*, *encryption*, *RADIUS*.

authorization—Permission for a user to carry out a certain set of tasks after he or she has access to your LAN. Authorization occurs *after* authentication is complete. On the MAX, you configure authorization to restrict access to:

- The terminal-server software.
- Simple Network Management Protocol (SNMP) manager utilities.
- Certain Domain Name System (DNS) servers.

See also *authentication.*

Auto-BERT—Automatic Byte-Error Test. During the Auto-BERT, the MAX monitors the entire data stream between codecs. At the end of the time period, if any channels have failed, the MAX clears the bad channels, redials, and repeats the test. The Call Status window displays BERT MAST at the dialing end of the call, and BERT SLAVE at the answering end of the call. The following status windows display the results of the Auto-BERT:

- The Line Errors window displays errors recorded on all current channels.
- The Session Errors window for a specific AIM port displays the cumulative error count for all channels connected to the port.
- The Port Info window displays the quality of all active calls.
- The Statistics window displays the quality of a call on a specific AIM port.

The maximum number of errors that can accumulate per channel is approximately 65,000. Note that the MAX reports the total number of errors for each channel during the current call, not the error rate. The MAX resets the error display for the current call to 0 (zero) when the call disconnects, or if the MAX disconnects a channel during the Auto-BERT or during the call itself. You can abort the Auto-BERT at any time by choosing the command DO Beg/End BERT.

Automatic Byte-Error Test—See *Auto-BERT*.

Automatic Number Identification—See *ANI*.

Autonomous System—See *AS*.

Autonomous System Border Router—See *ASBR*.

autosensing—A feature that enables you to change the device attached to an Ethernet port without reconfiguring the MAX.

Auxiliary Unit Interface—See *AUI*.

Average Line Utilization—See *ALU*.

B

B8ZS—Bipolar with 8-Zero Substitution. An encoding method in which alternating positive and negative voltage represents a 1, zero voltage represents a zero, and at least one bit out of every eight bits must be a 1.

backbone—The part of the communications network designed to carry the bulk of the traffic. The backbone provides connectivity between subnets in an enterprise-wide network. See also *enterprise-wide network, IP subnet*.

backbone area—An Open Shortest Path First (OSPF) area that connects routers for the purpose of hierarchical routing. The backbone area is special and always has the area number 0.0.0.0. To tie areas together, some routers belong to the backbone area and one other area. These routers are called *Area Border Routers (ABRs)*. See also *ABR*, *area*, *OSPF*, *router*.

backbone network—A network with a central cabling scheme linking it to other networks. Hosts on networks linked to the backbone can communicate with one another.

backbone router—A router attached to a backbone network by nailed-up lines. Usually, a backbone router does not have any built-in digital dial-up WAN interfaces. Manufacturers of backbone routers include Cisco, Wellfleet, 3Com, and CrossCom. See also *backbone network*, *router*.

backoff queue—A file in which the RADIUS accounting server stores unacknowledged records. See also *accounting server*, *RADIUS*.

back panel—The back portion of the MAX. The back panel includes slots, interfaces, and the following Ethernet interface LEDs:

LED	Description
ACT (Activity)	This LED is on when the MAX is detecting activity (network traffic) on its Ethernet interface.
COL (Collisions)	This LED is on when the MAX detects packet collisions on the Ethernet.
FDX	When this LED is on it indicates full duplex on the Ethernet.
100ST	When this LED is on, it indicates 100BT. When it is off, it indicates 10BT.
LINK (Link integrity)	This LED is on when the Ethernet interface is functional.

back panel alarm relay—See *alarm relay*.

back-to-back connection—A link in which the output of the sending device is connected directly to the input of the receiving device.

backup—The ability of the system to establish and use a temporary, alternate connection to a destination when the primary connection becomes unavailable. A backup connection replaces the primary connection, which must be a nailed-up (permanent) connection. The backup interface can be nailed-up or switched.

When the system detects that the primary interface is unavailable, it puts the primary interface in a Backup Active state. *It does not remove the routes to the primary interface*. It then diverts traffic from the primary to the backup interface. When the system detects that the primary interface is available again, it diverts traffic back to the primary interface. If the backup interface is a switched connection, the MAX then brings it down. One of the side effects of the datalink-layer backup interface is that, when a nailed-up interface specifies a backup interface, the routes to the nailed interface never go down.

You can specify a backup interface for a nailed-up connection in local Connection profiles, or in RADIUS. Nested backups are not supported. (The profile for a backup interface cannot specify another backup interface.) The profile for a backup interface does not inherit attributes, such as filters or firewalls, from the profile for the primary nailed-up connection.

See also *nailed-up circuit*, *switched circuit*.

backup and overflow—See *FT1-B&O*.

Backup Designated Router—See *BDR*.

Backward Explicit Congestion Notification—See *BECN*.

BACP—Bandwidth Allocation Control Protocol. BACP enables an Ascend unit to change bandwidth on demand by means of a standard set of rules, minimizing the need for the user to be involved in complex configuration issues. Especially useful in ISDN environments, BACP also enables the called device to call back the dial-in device over multiple links. The eight BACP sponsors are Ascend Communications, Bay Networks, Cisco Systems, Microsoft Corporation, Shiva Corporation, 3Com Corporation, U.S. Robotics and Xylogics. Detailed information on the current draft of the BACP specification is available through IETF lists.

bandwidth—The amount of data a link can carry, measured bits per second (bps) for digital signals, and in hertz (Hz) for analog signals. See also *analog signal*, *digital signal*.

Bandwidth Allocation Control Protocol—See *BACP*.

Bandwidth Allocation Protocol—See *BAP*.

bandwidth on demand—A MAX feature that reduces costs by automatically determining whether additional channels are required during periods of peak usage. The MAX releases the additional channels when they are no longer necessary.

Bandwidth ON Demand Interoperability Group—See *BONDING*.

banner—The text that first appears when a user logs into the terminal server. The default is `** Ascend MAX Terminal Server **`.

BAP—Bandwidth Allocation Protocol. BAP is a PPP protocol for managing bandwidth between two peers. Using BAP, the peers coordinate the process of adding and removing bandwidth. See also *BACP*, *PPP*.

base channel count—The number of channels to use for a Multilink Protocol (MP) connection. Because MP does not support Dynamic Bandwidth Allocation (DBA), the number of channels is fixed for the duration of the session. See also *DBA*, *MP*.

Basic Rate Interface line—See *ISDN BRI line*.

baud rate—The number of times a signal can switch from one state to another within one second. The more times a switch can occur, the higher the baud rate.

B channel—A 64-Kbps channel that carries user data. A B channel is a bearer channel, one of the fundamental components of the ISDN interface. See also *E1 PRI line*, *ISDN*, *ISDN BRI line*, *T1 PRI line*.

B-channel bundling—A technique for putting multiple voice conversations on a single line. Speech is divided so that bits are transmitted only when someone is speaking. In T1 multiplexing, bundles consist of four bits, represent 11 channels of 32-Kbps compressed data, and have an associated signaling Delta channel. See also *B channel*.

BDR—Backup Designated Router. A BDR is the router that the Open Shortest Path First (OSPF) area uses in the event that the Designated Router (DR) goes out of service. To prevent the DR from becoming a serious liability to the network if it fails, OSPF elects a Backup Designated Router (BDR). Other routers maintain adjacencies with both the DR and BDR, but the backup router leaves as many processing tasks as possible to the DR. If the DR fails, the backup immediately becomes the DR and a new backup is elected.

The MAX can function as either a DR or a BDR. However, many sites choose to assign LAN-based routers to these functions in order to dedicate the MAX to WAN processing. See also *adjacency*, *area*, *DR*, *OSPF*, *router*.

bearer channel—See *B channel*.

bearer service—An ISDN service for transmitting information from one device to another. Common bearer services are circuit-switched and Frame Relay services. See also *circuit switching*, *Frame Relay*.

BECN—Backward Explicit Congestion Notification. BECN is a bit set in a Frame Relay header to notify a source node that there is traffic congestion on the network. See also *FECN*, *Frame Relay*.

Bell 103—A carrier standard created by Bell Labs in the 1960s and 1970s. It accommodates modem-to-modem speeds of up to 300 bps, and is equivalent to the ITU-T V.21 standard. See also *ITU-T*, *V.21*.

Bell 212A—A carrier standard created by Bell Labs in the 1960s and 1970s. It accommodates modem-to-modem speeds of up to 1200 bps, and is equivalent to the ITU-T V.22 standard. See also *ITU-T*, *V.22*.

BER—Bit-Error Rate. The BER is the amount of received bits with errors as a percentage of the total number of bits received. It is commonly expressed as a number to the power of 10.

BERT—Bit-Error Rate Test. The BERT calculates the number of received bits with errors as a percentage of the total number of bits received. The Bit-Error Rate (BER) is commonly expressed as a number to the power of 10.

BGP—Border Gateway Protocol version 4. BGP routes packets between networks that use different types of protocols. See also *EGP*.

binary data—Data in the form of zeroes and ones.

Binary mode—The Telnet 8-bit Binary option. You can run X -Modem and other 8-bit file transfer protocols using this mode. In 8-bit Binary mode, the Telnet escape sequence does not operate. The Telnet session can close only if one end of the connection quits the session. If you are a local user not connected through a digital modem, the remote-end user must quit. A user can override the Binary setting on the Telnet command line. Compare with *ASCII mode*, *Transparent mode*. See also *Telnet*, *Telnet mode*.

Binary Local mode—A data-transfer mode for X.25/T3POS calls. Binary Local mode specifies that there is no error recovery between the T3POS Packet Assembler/Disassembler (PAD) and the host, but that error recovery is in place between the PAD and the DTE. Like Blind mode, it passes data between the DTE and the host without reference to the protocol being used. Unlike Blind mode, Binary Local specifies that the system continues to use the T3POS protocol between the DTE and the PAD. Compare with *Blind mode*, *Local mode*, *Transparent mode*. See also *DTE*, *PAD*, *X.25/T3POS*.

Bipolar with 8-Zero Substitution—See *B8ZS*.

B-ISDN—Broadband-Integrated Services Digital Network. B-ISDN is a very high-speed data service, providing data transmission at rates higher than T1 or E1. See also *broadband*, *E1 line*, *ISDN*, *T1 line*.

Bit—Binary digit, the smallest unit of information a computer can process, representing one of two states (indicated by 1 and 0).

bit inversion—A method of turning data 1s into 0s and data 0s into 1s. Bit inversion applies only to calls between codecs. In some connections, you need to invert the data to avoid transmitting a pattern that the connection cannot handle. If you apply bit inversion, you should do so on both sides of the connection.

bit rate—The number of bits that travel over a connection per second. See also *bps*.

bits per second—See *bps*.

black-hole interface—An interface that enables the router to handle packets whose IP address matches an unused IP address in a summarized address pool. The black-hole interface has an IP address of 127.0.0.3. When you specify this address as the router to the destination pool network, the MAX silently discards packets to an invalid host on that network. See also *pool summary*.

Blind mode—A data-transfer mode for X.25/T3POS calls. Blind mode specifies that the T3POS Packet Assembler/Disassembler (PAD) does not provide any error recovery. In this mode, the DTE and the host system provide error recovery for the connection. In addition, the T3POS PAD does not clear a call when it receives a clear-request command from the DTE. The PAD or the host system must clear the call. Finally, the PAD passes all data *blindly,* without regard to the protocol in use. This mode provides a method of passing raw binary data between the DTE and the host system without reference to the protocol being used. Compare with *Binary Local mode*, *Local mode*, *Transparent mode*. See also *DTE*, *PAD*, *X.25/T3POS*.

BONDING—Bandwidth ON Demand Interoperability Group. BONDING is a consortium of over 40 data-communications-equipment vendors and service providers who joined together to create a standardized inverse-multiplexing protocol. The BONDING protocol enables inverse multiplexers from different vendors to interoperate. BONDING also refers to the resultant specification, sometimes known as the *BONDING specification*. See also *inverse multiplexer*, *inverse multiplexing*.

BOOTP—Bootstrap Protocol. BOOTP enables a network user with a diskless workstation to receive an IP address and have the network operating system automatically started on the workstation. A BOOTP server assigns the IP address from a pool of addresses for a fixed duration.

The MAX can use BOOTP to get settings and check for a new software load. In addition, you can enable the terminal server to respond to BOOTP within a Serial Line Internet Protocol (SLIP) session. An interactive user who initiates a SLIP session can get an IP address from a designated IP address pool by means of BOOTP.

BOOTP is the basis for Dynamic Host Configuration Protocol (DHCP), a more advanced protocol.

See also *BOOTP relay*, *BOOTP request*, *BOOTP server*, *DHCP*, *IP address*, *IP address pool*, *SLIP*, *terminal server*.

BOOTP relay—A method of sending (*relaying*) Bootstrap Protocol (BOOTP) requests to other networks. On the MAX, you specify the IP address of a BOOTP server for handling BOOTP requests. If a server is on the same LAN as the MAX, BOOTP requests from other networks are relayed to the server. If a server is on another network, BOOTP requests from clients on the same LAN as the MAX are relayed to the remote server. If you specify two BOOTP servers, the MAX that relays the BOOTP request determines when each server is used. See also *BOOTP*, *BOOTP request*, *BOOTP server*.

BOOTP request—A request a client makes to a BOOTP server in order to receive an IP address or start the operating system of a network workstation. See also *BOOTP*, *BOOTP relay*, *BOOTP server*.

BOOTP server—A server that handles BOOTP requests from network clients. See also *BOOTP*, *BOOTP relay*, *BOOTP request*.

Bootstrap Protocol—See *BOOTP*.

Border Gateway Protocol version 4—See *BGP*.

bps—A nested acronym, meaning binary digits per second, and a measure of the capacity of a line.

bridge—A hardware device that transmits packets between networks. A bridge forwards packets from one network to another, and discards packets destined for hosts on the sending network. Operating at the Data Link layer, a bridge makes multiple networks look like a single network to higher-level protocols and software. See also *Data Link layer*.

bridge entry—An entry in a bridging table. The MAX is a transparent bridge (also called a *learning bridge*). As the MAX forwards a packet, it notes the packet's source address and creates a bridging entry that associates a host's Media Access Control (MAC) address with a particular Ethernet interface.

The MAX also learns about bridging links from Connection profiles and RADIUS user profiles, either because the remote caller used the profile to dial the link, or because the profile matched an incoming call. In addition, you can specify static bridge entries in a local profile or RADIUS pseudo-user profile. See also *bridge*, *bridging table*, *Connection profile*, *Ethernet*, *MAC*, *pseudo-user profile*, *user profile*.

bridging—A method of moving packets between networks by means of a device called a *bridge*, which operates at the Data Link layer. In Figure 7, the MAX at site A acts as a bridge between the Ethernet and site B.

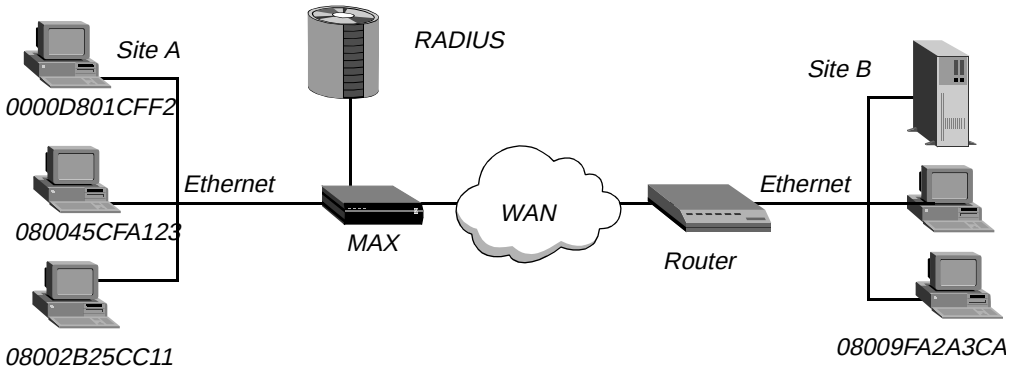

Figure 7. Bridging configuration

The MAX at site A gradually learns addresses on both networks by looking at each packet's source address, and it develops a bridging table that includes the following entries:

```
0000D801CFF2        SITEA
080045CFA123        SITEA
08002B25CC11        SITEA
08009FA2A3CA        SITEB
```

If the MAX receives a packet whose destination MAC address is not on the local network, it first checks its internal bridging table. If it find the packet's destination MAC address, the MAX dials the connection and bridges the packet. If it does not find the address, the MAX checks for active sessions that have bridging enabled. If one or more active bridging links are up, the MAX forwards the packet across all active sessions that have bridging enabled.

See also *bridge*, *bridge entry*, *bridging table*, *Data Link layer*.

bridging table—A table that contains entries pairing up a host's Media Access Control (MAC) address with a particular Ethernet interface. If the MAX receives a packet whose destination MAC address is not on the local network, it first checks its bridging table. If it find the packet's destination MAC address, the MAX dials the connection and bridges the packet. If it does not find the address, the MAX checks for active sessions that have bridging enabled. If one or more active bridging links are up, the MAX forwards the packet across all active sessions that have bridging enabled. See also *bridge*, *bridge entry*, *Ethernet*, *MAC*.

BRI line—See *ISDN BRI line*.

BRI/LT slot card—See *IDSL card*.

broadband—A data communications technology that transmits data in channels and uses those channels simultaneously.

Broadband Integrated Services Digital Network—See *B-ISDN*.

Broadband ISDN—See *B-ISDN*.

broadcast—The process of sending a message to all connected hosts, as opposed to sending a message to a single host or to members of a multicast group. See also *broadcast network*, *broadcast packet*, *multicast*, *multicast group*.

broadcast network—A network in which the router sends packets to all users, whether they appear on subscription lists or not. In an Open Shortest Path First (OSPF) topology, a broadcast network is any network that has more than two OSPF routers attached and can address a single physical message to all of them. See also *OSPF*, *multicast network*, *router*, *unicast network*.

broadcast packet—A packet containing a broadcast address, which indicates that all connected hosts receive the message. See also *broadcast*, *broadcast network.*

build—The name of the software binary you install on the MAX. For example, `ti.m40` is the MAX 4000 T1 IP-only software build. For the names of all the software builds and the features they provide see `/pub/Software-Releases/Max/Upgrade-Filenames.txt` or `/pub/Software-Releases/Pipeline/Upgrade-Filenames.txt` on the Ascend FTP server.

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

builddbm file—A file that enables you to create a UNIX DBM database for use with the Ascend RADIUS daemon. See also *DBM database*, *RADIUS daemon.*

buildout—For a T1 line with an internal Channel Service Unit (CSU), the amount of attenuation the MAX should apply to the line's network interface in order to match the cable length from the MAX to the next repeater. When you specify a buildout value, the MAX applies an attenuator to the T1 line, causing the line to lose power when the received signal is too strong. Repeaters boost the signal on a T1 line. If the MAX is too close to a repeater, you need to add some attenuation. See also *attenuation*, *attenuator*, *CSU*, *repeater*, *T1 line*.

bundle—A group of physical links, such as multiple asynchronous lines, or multiplexed links, such as Multilink Protocol (MP), Multilink Protocol Plus (MP+), X.25, or Frame Relay connections. The links in a bundle can be of different types, such as dial-up asynchronous and nailed-up synchronous connections. See also *bundle owner*.

bundle owner—The MAX that answers the first call in the MP or MP+ bundle. If a bundle spans more than one MAX in a stack, an exchange of information flows between the MAX units in the bundle. If the call belongs to an existing bundle, the MAX that answered and the bundle owner exchange information about the bundle. Furthermore, the MAX that answered the call forwards all incoming data packets over the Ethernet to the bundle owner.

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

See also See also *bundle*, *MP*, *MP+*, *stacks*.

bus—A path for signals transmitted between a computer's CPU and other hardware devices.

byte—8 bits of data, also called an *octet*.

byte offset—See *offset*.

C

cached token—A password dynamically generated on a token card, transmitted by Challenge Handshake Authentication Protocol (CHAP), and then cached for reuse. When the MAX needs to add channels or make a new call, the MAX uses the cached token to authenticate the additional bandwidth. You can specify a timeout value for the cached token, or configure the system to maintain the token throughout the session. See also *Cache-Token authentication, CHAP, token, token card, token-card authentication, token-card server*.

Cache-Token authentication—An authentication method that uses Challenge Handshake Authentication Protocol (CHAP) to transmit the initial token, and then caches the token for reuse. The system later uses the cached token when the MAX adds new channels or makes a new call. See also *ACE authentication, cached token, SafeWord authentication, token, token card, token-card authentication, token-card server.*

cable modem—In the early stages of testing, a device that promises to deliver high-speed data throughput over the coaxial cables used by the cable TV industry. Compare with *digital modem, modem*.

call—A single session in which a calling device and an answering device connect over the WAN.

callback—A type of security in which you instruct the MAX to hang up and call back when it receives an incoming call. You can require callback to ensure that the MAX makes a connection with a known device. Hanging up and calling back adds a level of certainty that the connection is with a trusted user, especially because the MAX calls back immediately after verifying the user's name and password. For the MAX to use callback, it must be able to both receive and initiate calls. Callback security applies only to switched lines. See also *authentication, switched line*.

Callback Control Protocol—See *CBCP*.

call blocking—A MAX feature that enables you automatically stop the unit from attempting to place an outgoing call on a connection that repeatedly fails. Successive retries can cause excessive charges, congestion, and performance problems. Using call blocking, you can prohibit additional retries after a specified number of failed connection attempts. You can also control the length of time call blocking is in effect.

Call-Connected packet—A packet sent by remote Data Terminal Equipment (DTE) when the device accepts a call from a unit on an X.25 network. See also *DTE*, *X.25*.

Call Detail Reporting—See *CDR.*

called-number authentication—A form of authentication in which the MAX uses the called-party number to authenticate the connection. The remote end uses this form of authentication to make sure that the call goes to a known destination. When the profile requires called-number authentication, the number called must match a phone number in a Connection profile or RADIUS user profile. The MAX also uses the called number to direct incoming calls to a particular device. See also *called-party number, Connection profile, user profile.*

called-party number—An information element of the Q.931 ISDN signaling protocol. The called-party number is the phone number the remote device calls to connect to the MAX, but without a trunk group or dialing prefix specification. This number is always available if specified in a profile. See also *called-number authentication*.

caller ID—See *CLID*.

Caller Identification—An ISDN telephone service that enables the called party's equipment to display the telephone number of the caller. See also *Caller Identification Restriction.*

Caller Identification Restriction—An ISDN telephone service that enables the caller to prevent his or her telephone number from being displayed to the called party. See also *Caller Identification*.

call filter—A packet filter that defines which packets can bring up a connection or reset the idle timer for an established link (Figure 8). A call filter prevents unnecessary connections and helps the MAX distinguish active traffic from "noise."

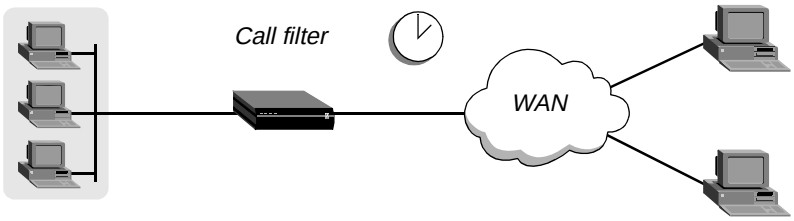

Figure 8. Call filters can prevent certain packets from resetting the timer

By default, any traffic to a remote site triggers a call, and any traffic across an active connection resets the connection's idle timer. When a session's idle timer expires, the MAX terminates the session. The idle timer is set to 120 seconds by default, so if a connection is inactive for two minutes, the MAX terminates the connection.

When you apply a call filter, its forwarding action does not affect which packets are sent across an active connection. For a call filter to prevent an interface from remaining active unnecessarily, you must define rules for both incoming and outgoing packets. Otherwise, if only input rules are defined, outgoing packets keep a connection active. If only output rules are defined, incoming packets keep a connection active.

Compare with *data filter*. See also *generic filter, input filter, IP filter, IPX filter, output filter, packet filter*.

Calling-Line ID—See *CLID.*

Calling-Line ID authentication—See *CLID authentication*.

call logging—A method of logging call information from the MAX. Based on RADIUS accounting, call logging enables you to keep records for resource management or troubleshooting. When you set up call logging, you can create duplicate accounting information for sites that wish to keep accounting records separate from call logging records. The MAX sends Start session, Stop session, and Failure-to-start session packets to a call-log host. The call-log information is sent independently of RADIUS accounting records. If both call logging and RADIUS accounting are in use, the information is sent in parallel. See also *accounting*, *call-log host*, *Failure-to-start session*, *Start session*, *Stop session*.

call-log host—A local host that supports the RADIUS accounting protocol and is configured to communicate with the MAX. See also *accounting*, *call logging*.

call request—On an X.25 network, a request made by the calling party, asking the Data Terminal Equipment (DTE) to accept the call. See also *Call-Request packet*, *Call-Request timer*, *DTE*, *X.25*.

Call-Request packet—A packet sent by a device when it makes an outgoing call on an X.25 network. See also *call request*, *Call-Request timer*, *X.25*.

Call-Request timer—The number of ten-second ticks the MAX waits before clearing an outgoing call that the remote Data Terminal Equipment (DTE) has not accepted. When a device makes an outgoing call, it sends a Call-Request packet. If the remote DTE accepts the call, it sends back a Call-Connected Packet. If the DTE refuses the call, it sends back a Clear-Request packet. See also *Call-Connected packet*, *call request*, *Call-Request packet*, *Clear-Request packet*, *DTE*, *X.25*.

call routing—The way that the MAX routes incoming and outgoing calls to the proper modules.

When the MAX receives a call on a WAN line, it performs Calling-Line ID (CLID) or called-number authentication (if appropriate), answers the call, and determines which slot should receive the call. It then finds the caller's profile, authenticates the call, builds a session, and passes the data stream to the appropriate module or host. When a call is routed to the Ethernet port, the bridge/router software forwards it to a host according to the packet's destination address.

When the MAX dials an outgoing call, it routes the call from the originating slot to a WAN channel. It first looks for channels associated with the trunk group specified by the dialed number (if any) and the port that originated the call. If no trunks have available channels, the MAX does not place the call.

See also *called-number authentication*, *CLID authentication.*

Call Setup message—See *ISDN Call Setup message*.

call spanning—A method of enabling a Multilink Protocol (MP) or Multilink Protocol Plus (MP+) call to span the MAX units in the stack. Call spanning using a stack configuration can be effective when:

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

Call spanning is protocol independent, and should work with all protocols supported by the MAX. See also *hunt group*, *MP*, *MP+*, *rotary hunt group*, *stacks*.

Call User Data—See *CUD*.

CAP—Competitive Access Provider. A CAP is a business that competes with the local telephone company in providing clients with access to services. For example, a cable company that offers high-speed data communications services is a CAP.

Carrier Detect—See *CD*.

carrier services—Telecommunications services provided to the public for a fee, such as ISDN lines and Frame Relay services.

CAS—A carrier switch type in New Zealand.

cause code—A numerical diagnostic code sent from an ISDN switch to Data Terminal Equipment (DTE). A cause code indicates why call-establishment failed, or why a call was terminated. The cause codes are part of ISDN D-channel signaling communications supported by the Signaling System 7 supervisory network. When you dial a call from the MAX using ISDN access, the MAX reports the cause codes in the Message Log status menu. When the MAX clears the call, it reports a cause code, even when inband signaling is in use. A cause code is also called a *cause element*. See also *DTE, ISDN, Signaling System 7*.

cause element—See *cause code*.

cause field—A field that indicates an event that triggered an X.25 Clear-Request, Reset-Request, or Restart-Request packet. Values for the cause field can vary, depending on the packet type. See also *Clear-Request packet, diagnostic field, Reset-Request packet, Restart-Request packet, X.25*.

CBCP—Callback Control Protocol. Microsoft's CBCP is a Link Control Protocol (LCP) option negotiated at the beginning of Point-to-Point Protocol (PPP) sessions. CBCP authenticates a caller by means of a user name and password, and offers additional security.

Microsoft developed CBCP to address a need for greater security when establishing PPP connections. The standard callback option defined in RFC 1570 has a potential security risk because the authentication is performed after the callback. CBCP callback, like Ascend's proprietary callback, occurs after authentication, leaving no potential security hole. CBCP also offers features not available with standard callback. The client side supports a configurable time delay, enabling users to initialize modems or startup software before the MAX calls the client. You can configure the MAX with a phone number to use for the callback, or you can allow the client to specify the phone number.

Currently, Microsoft's Windows NT 4.0 and Windows 95 software support client-side authentication using CBCP. The MAX now supports a CBCP central-site solution. While support for CBCP is configured systemwide on the MAX, not every connection must negotiate its use. The calling and called sides of a PPP session initiate authentication after acknowledging that CBCP is to be used. Currently, the MAX does not initiate LCP negotiation of CBCP. The MAX responds to *caller* requests to configure CBCP.

See also *callback, LCP, PPP*.

C-bit Parity Errors—See *CPERR*.

CCITT—Consultative Committee on International Telegraphy and Telephony. The CCITT is a disbanded organization whose standards were moved to the UN-sanctioned ITU-T on March 1, 1993. See also *ITU-T*.

CCP—Compression Control Protocol. CCP enables both ends of a Point-to-Point Protocol (PPP) connection to negotiate whether to use data compression, and if so, which algorithm to use. See also *data compression*.

CD—Carrier Detect. CD is a signal sent from a modem to a host, indicating that the modem is online. CD can also be referred to as *Data Carrier Detect (DCD)*.

CDMA—Code Division Multiple Access. CDMA is a digital wireless transmission technique that uses mathematical codes, instead of frequencies or time slots, to transmit information. CDMA is a leading digital standard. See also *AMPS*, *CDPD*, *cellular communication*, *cellular modem*, *cellular network*, *wireless technology*.

CDPD—Cellular Digital Packet Data. CDPD is a digital wireless transmission technique that uses idle voice channels on the existing Advanced Mobile Phone Service (AMPS) cellular telephone network. CDPD transmits data packets at a raw data rate of 19.2 Kbps, using channel hopping to move data packets through unused spaces across different frequencies. Because data is not a time-sensitive a service as voice, data can be fragmented and then reassembled at the receiving end. CDPD is particularly suited to sending small messages and transactions. It is not appropriate for transmitting multimegabit files. See also *AMPS*, *CDMA*, *cellular communication*, *cellular modem*, *cellular network*, *wireless technology*.

CDR—Call Detail Reporting. CDR is a feature that provides a database of information about each call, including date, time, duration, called number, calling number, call direction, service type, and associated inverse multiplexing session and port. Because the network carrier charges for bandwidth on an as-used basis, and bills each connection in an inverse-multiplexed call as a separate charge, you can use CDR to understand and manage bandwidth usage and the cost of each inverse-multiplexed session.

cell relay—See *ATM*.

cellular communication—A type of wireless communication first available in 1981. To implement this technology, a carrier divides a city or county into units called *cells*. Each cell contains the transmitters and receivers that provide the telephone service. The frequencies assigned to one cell are limited to the boundaries of that cell. When a cellular phone moves from one cell toward another, a computer at the switch monitors the motion and hands off the phone call to the new cell, which uses another radio frequency. The transfer is not noticeable to the user. Compare with *landline telephone communication*. See also *CDMA*, *CDPD*, *cellular modem*, *cellular network*, *wireless technology*.

Cellular Digital Packet Data—See *CDPD*.

cellular modem—A modem that transmits data between remote locations using cellular technology. See also *CDMA*, *CDPD*, *cellular communication*, *cellular network.*

cellular network—A network that enables cellular subscribers to travel anywhere in the country and remain connected to the Public Switched Telephone Network (PSTN) by means of their mobile phones. See also *CDMA*, *CDPD*, *cellular communication*, *cellular modem.*

Central Office—See *CO*.

Central Processing Unit—See *CPU*.

Centrex—Business telephone service that a Local Exchange Carrier (LEC) offers from a central office. Centrex services include call forwarding, call transfer, call restrictions, and call hold. Centrex service is an alternative to buying or leasing a Private Branch Exchange (PBX). See also *LEC*, *PBX*.

Challenge Handshake Authentication Protocol—See *CHAP*.

Change-Filter-Request packet—A request to change the packet filters for a bridging/routing session. See also *Change-Filter-Request-ACKed packet*, *Change-Filter-Request-NAKed packet*.

Change-Filter-Request-ACKed packet—A message the MAX sends if it found at least one bridging/routing session for which it could change packet filters. Compare with *Change-Filter-Request-NAKed packet*. See also *Change-Filter-Request packet*.

Change-Filter-Request-NAKed packet—A message the MAX sends if it could not find a bridging/routing session for which it could change packet filters. Compare with *Change-Filter-Request-ACKed packet*. See also *Change-Filter-Request packet*.

channel—A portion of a line's bandwidth. A line contains a fixed number of channels. Each line can contain switched channels only, nailed-up channels only, or a combination of switched and nailed-up channels. See also *bandwidth*, *line*, *nailed-up channel*, *switched channel*.

channelized T1 PRI/E1 PRI—A T1 PRI or E1 PRI line divided into individual 64-Kbps channels, or into channels whose data rate is a multiple of 64 Kbps (such as a 256-Kbps channel made from four 64-Kbps channels). Channelized T1 PRI or E1 PRI lines can consist of switched lines with inband signaling, or nailed-up lines. For example, a nailed-up line can run from the Central Office (CO) to the corporate headquarters as a single, unchannelized T1 PRI or E1 PRI line, and can then be divided into channels when it runs to remote sites from the corporate headquarters. See also *E1 line*, *E1 PRI line*, *inband signaling*, *nailed-up line*, *switched line*, *T1 line*, *T1 PRI line*, *unchannelized service*.

Channel Service Unit—See *CSU*.

CHAP—Challenge Handshake Authentication Protocol. CHAP authentication verifies the caller's identity by using a three-way handshake upon initial link establishment, and then by repeating the handshake any number of times. In CHAP authentication, the authentication server sends a challenge to the caller. The caller responds with an MD5 digest calculated from the password. The authentication server then checks the digest against its own calculation of the expected hash value to authenticate the call. The server can send a new challenge at random intervals.

CHAP is a stronger authentication method than Password Authentication Protocol (PAP), because the password does not travel across the line as plain text. In addition, the use of repeated challenges limits the time of exposure to any single attempt to break the encryption code, and the server is in control of how often it sends challenges. See also *encryption*, *PAP*.

Char-to-Char timer—For an X.25/T3POS connection, a timer that indicates the maximum amount of time permitted between characters sent from the Data Terminal Equipment (DTE) to the Packet Assembler/Disassembler (PAD). The Char-to-Char timer is also called the *T1 timer.* See also *DTE*, *PAD*, *X.25/T3POS*.

Checkpoint record—A RADIUS accounting record that enables you to retrieve information on each user session in the event of network disruption. By default, RADIUS Accounting logs a Start and Stop record for each user session. If a disruption in service causes a connection to go down before the MAX receives a RADIUS Stop record, you can use the Checkpoint records to reconstruct usage.

In the RADIUS detail file, a Checkpoint record contains the same group of attributes as a RADIUS Stop record. However, the value for the Acct-Status-Type attribute in a Checkpoint record is the number 3. When queuing RADIUS Accounting records, the MAX prioritizes Start and Stop records ahead of Checkpoint records.

See also *accounting*, *Start record*, *Stop record*.

circuit—In general, a connection between endpoints over a physical medium. On a Frame Relay network, a circuit is a Permanent Virtual Circuit (PVC) segment that consists of two Data Link Connection Indicator (DLCI) endpoints and possibly two Frame Relay profiles. It requires two and only two DLCI numbers. Data is dropped if the circuit has only one DLCI. If more than two are defined, only two are used. Circuits are defined in two Connection profiles or two RADIUS user profiles. Data coming in on the DLCI configured in the first profile is switched to the DLCI configured in the second one. See also *DLCI*, *Frame Relay*, *PVC*.

circuit connection—A connection that follows a specified path through the Frame Relay switch. By linking two Data Link Connection Indicator (DLCI) endpoints, the MAX creates a Permanent Virtual Circuit (PVC). The two DLCI endpoints act as a tunnel. A circuit connection is illustrated in Figure 9.

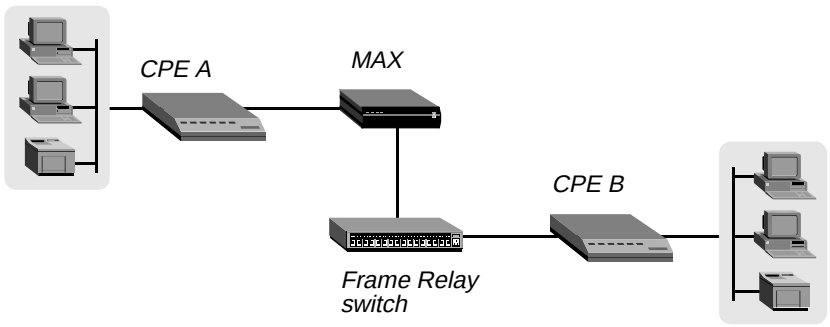

Figure 9. Circuit connection

Data that the MAX receives on one DLCI bypasses the Ascend router and goes out on the other DLCI. If any one of the DLCIs in a PVC becomes inactive because of a disconnect or failure, the PVC using that DLCI becomes inactive. A physical line can carry multiple DLCIs, and the failure of the line causes the failure of all the DLCIs it carries. Compare with *Frame Relay Direct*, *Frame Relay gateway*. See also *DLCI*, *Frame Relay switch*, *PVC*, *router*.

circuit-level inverse multiplexing—A method of inverse multiplexing in which the inverse multiplexer slices the data stream into equal portions, and transmits each portion over an available circuit. The receiving end adjusts for network-induced delay and reassembles the data packets into their proper order. The AIM and BONDING protocols define how circuit-level inverse multiplexing works. Applications that require transparent digital circuits (such as videoconferencing, nailed-up backup and overflow, and bulk file transfer) use circuit-level multiplexing. Compare with *packet-level inverse multiplexing*. See also *AIM*, *BONDING*, *inverse multiplexer*, *inverse multiplexing*.

Circuit-Switched Cellular Data—See *CSCD*.

circuit-switched line—A temporary connection, like a telephone call. A temporary connection can be made to a variety of sites to handle occasional data-transfer needs or to provide additional bandwidth when needed.

Circuit-Switched Public Data Network—See CSPDN.

circuit switching—A mode of data transfer in which a dedicated connection is busy for the duration of the call. Compare with *packet switching*.

Clear-Confirmation packet—On an X.25 network, a packet that the Data Terminal Equipment (DTE) or Data Circuit-Terminating Equipment (DCE) receives in response to its request to clear a call. When the device receives a Clear-Confirmation packet from the remote end, the call is cleared and the logical channel is available for other calls. See also *Clear-Indication packet*, *Clear-Request packet*, *DCE*, *DTE*, *X.25*.

Clear-Indication packet—On an X.25 network, a packet that the Data Circuit-Terminating Equipment (DCE) sends to refuse an incoming call, or to clear a call when the data exchange is complete. See also *Clear-Confirmation packet*, *Clear-Request packet*, *DCE*, *X.25*.

Clear-Request packet—On an X.25 network, a packet that the Data Terminal Equipment (DTE) sends to refuse an incoming call, or to clear a call when the data exchange is complete. See also *Clear-Confirmation packet*, *Clear-Indication packet*, *DTE*, *X.25*.

Clear-Request retries—The number of times the MAX sends a Clear-Request packet on an X.25 network before waiting indefinitely for a response. See also *Clear-Request packet*, *X.25*.

Clear-Request timer—The number of ten-second ticks that the MAX waits before retransmitting a Clear-Request packet. See also *Clear-Request packet*, *X.25*.

Clear To Send—See *CTS*.

CLID—Calling-Line ID. The CLID is the telco-provided phone number of the calling device that wants to connect to the MAX. A CLID is also known as a *caller ID*. See also *CLID authentication*.

CLID authentication—Calling-Line ID authentication. CLID authentication is a method the MAX uses to authenticate incoming calls by checking the calling party's phone number as received from the telco. The CLID is the phone number of the calling device. The MAX performs CLID authentication before enabling the MAX to answer an incoming call. When the profile requires CLID authentication, the caller's phone number must match a phone number specified in a local Connection profile or RADIUS user profile. You can thereby ensure that the call comes from a known source.

You can use CLID authentication only where the call information is available end-to-end and Automatic Number Identification (ANI) applies to the call. In some areas, the WAN provider might not be able to deliver CLIDs, or a caller might keep a CLID private. Typically, a site uses CLID authentication to protect against a situation in which an unauthorized user obtains the name, password, and IP address of an authorized user, and then calls the MAX from another location.

See also *CLID*, *Connection profile*, *RADIUS*, *user profile*.

client—A user or device that requires services from another unit or program. For example, a user requesting access is a client of the MAX, and a MAX making a RADIUS authentication request is a client of the RADIUS server. See also *RADIUS server*.

client DNS—A configuration that enables the MAX to direct incoming connections to a Domain Name System (DNS) server belonging to a particular client or location, thereby preventing WAN users from accessing a local DNS server. Client DNS has two levels: a global configuration that applies to all PPP connections, and a connection-specific configuration. The MAX uses the global client addresses only if none are specified in the Connection profile. The addresses configured for client DNS servers are presented to WAN connections during IPCP negotiation. You can also choose to present your local DNS servers if no client servers are defined or available. See also *DNS*, *IPCP*.

clients file—A file that defines the client machines permitted to make requests to the RADIUS server. For the RADIUS daemon to respond to client requests from the MAX, you must enter a line specifying the MAX unit's name and password in the `clients` file. See also *RADIUS daemon*, *RADIUS server*.

clock—A timing mechanism for synchronizing data communication and processing tasks. A clock divides time into very short intervals. The clock speed is the number of intervals per second. See also *clock source*.

clock source—The master source for clocking of synchronous connections. The MAX uses a single synchronous clock source. The MAX chooses the clock source from the T1 or E1 lines you specify as possible external sources. If there are no eligible external sources, the system uses an internal clock. See also *clock*, *synchronous transmission*.

Closed User Group—See *CUG*.

CO—Central Office. The CO is the telephone switching office to which a customer directly connects. It connects the customer to other portions of the telephone network.

codec—COder/DECoder. A codec is a device that encodes analog data into a digital signal for transmission over a digital medium. Codecs are often used for videoconferencing. See also *analog data*, *digital signal*.

Code Division Multiple Access—See *CDMA*.

coldstart notification—In a RADIUS Accounting Stop record, a value that informs the accounting server that the MAX has started up. See also *accounting*, *accounting server*, *Stop record*.

command mode—A terminal-server mode in which you can enter commands at the terminal-server prompt. Compare with *immediate mode*, *menu mode*.

comment line—A line in a RADIUS user profile or pseudo-user profile that describes the purpose of one or more lines of the profile. Beginning with the # character at column one, the comment line consists of text that extends to the end of the line. You can embed a comment line anywhere in a profile. See also *pseudo-user profile*, *user profile*.

community name—A password that the MAX sends to the Simple Network Management Protocol (SNMP) manager when an SNMP trap event occurs, and that the manager sends to the MAX with each polling request. The password authenticates the sender. The default is `public`. See also *agent*, *manager*, *SNMP*.

Competitive Access Provider—See *CAP*.

Compressed Serial Line Internet Protocol—See *CSLIP*.

compression—A process that reduces the quantity of bandwidth or storage space required to encode a block of information. See also *data compression*, *link compression*, *slot compression*, *Stac compression*, *Stac-9 compression*, *Stac LZS compression*, *V.42bis*, *VJ compression*.

Connection profile—A local profile containing authentication and configuration information about a remote device or user.

Console session—An interactive configuration and management session established by means of a serial connection or Telnet link to the MAX.

Consultative Committee on International Telegraphy and Telephony—See *CCITT*.

Control frame—On an X.25/T3POS network, a supervisory frame of the format:

`<SOH MSS CUD STX [`*`data`*`] ETX XRC>`

where:

- `SOH` is the ASCII character \001.
- `MSS` is the Mode Selection Signal that can be (optionally) used to indicate the call mode.
- `CUD` is the Called User Data, which may contain an X.121 address in addition to user facilities or call user data in an X.28 format.
- *`data`* is optional in the control frame. In Transparent and Blind modes, the T3POS PAD is restricted to passing data frames between the T3POS Data Terminal Equipment (DTE) and the T3POS host.
- `ETX` is the ASCII character \003.
- `XRC` is the checksum. For all modes except Binary Local, the checksum is a one-character Longitudinal Redundancy Check (LRC) checksum. For Binary Local mode, the checksum is a two-character Cyclic Redundancy Check (CRC) checksum.

Control frames are in use only when a call is being established, and not during data transfer. See also *Binary Local mode*, *Blind mode*, *CRC*, *DTE*, *Local mode*, *Transparent mode*, *X.25/ T3POS*.

control-lead signaling—A method of toggling one or more leads within the cable between the application and the MAX in order to initiate a dialed call.

control-line state—A state that results when a device sends a signal through a pin and over the line to another device. The signal being sent determines the control-line state. For example, a device can send a signal to inform another party that it is ready to receive data. In this case, the control-line state is Data Transmit Ready (DTR). The process of sending control signals is called *handshaking*. See also *DTR*, *handshaking*.

Control port—A port on the MAX that connects to a VT100 terminal or modem to provide the menu-driven user interface. The Control port runs at 9600 bps, eight bits per character, no parity, no flow control, and one stop bit.

convergence—The time it takes all routers to receive information about a change to the network topology. A slow convergence can result in routing loops and errors. A Routing Information Protocol (RIP) router broadcasts its entire routing table every 30 seconds. On a 15-hop network, convergence can be as high as 7.5 minutes. In contrast, Open Shortest Path First (OSPF) uses a link-state database of the network, and propagates only changes to the database, resulting in faster convergence. See also *link-state database*, *OSPF*, *RIP*.

Coordinated Universal Time—See *UTC*.

cost—An Open Shortest Path First (OSPF) value you assign to the output side of each router interface. The cost indicates the likelihood that the MAX will use the interface to transmit data. The lower the cost, the more likely that the MAX will use the interface.

Figure 10 shows how costs are used to direct traffic over high-speed links. For example, if Router-2 in Figure 10 receives packets destined for Host B, it will route them through Router-1 across two T1 links (Cost=20) rather than across one 56kbps B-channel to Router-3 (Cost=240).

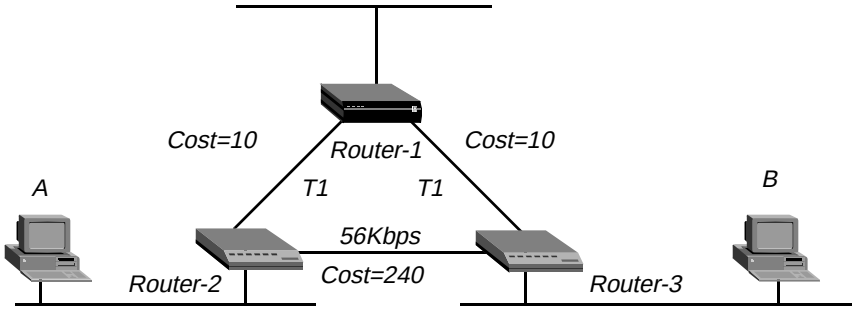

Figure 10. OSPF costs for different types of links

You can use the cost to perform preferred path selection. If two paths to a destination have equal costs, you can assign a higher cost to one of the paths, making it a backup when the primary path is not available. In addition, you may want to reflect the bandwidth of a connection when assigning costs. As in Figure 10, the cost of a single B-channel connection could be 24 times greater than the cost of a T1 link.

The MAX has a default cost of 1 for a connected route (Ethernet) and 10 for a WAN link. Be careful when assigning costs. Incorrect cost metrics can cause delays and congestion on the network.

See also *OSPF*, *route*, *router*.

CPE—Customer Premises Equipment. CPE is equipment connected to the telephone network, and located at the customer's site. The equipment can be owned or leased.

CPERR—C-bit Parity Errors. CPERR indicates the number of times that the C-bit parity check failed on the E1 line. See also *E1 line*.

CPU—Central Processing Unit. The CPU is a device's main processor.

CRC—Cyclic Redundancy Check. CRC is an error-detection method that uses a mathematical divisor to check the integrity of the data in a transmitted packet. Compare with *LRC*.

crossover cable—A cable with wires that cross over, so that the terminating ends of the cable have opposite wire assignments. Compare with *straight-through cable*.

CSCD—Circuit-Switched Cellular Data. CSCD is a wireless transmission technology that supports sending large files and faxes. CSCD uses switches to set up connections in analog cellular networks, and is also used in conjunction with such digital packet technologies as Cellular Digital Packet Data (CDPD). See also *CDPD*, *cellular communication*, *wireless technology*.

CSLIP—Compressed Serial Line Internet Protocol. CSLIP is a form of the Serial Line Internet Protocol (SLIP). Both SLIP and CSLIP enable you to transmit IP packets over serial connections, but CSLIP uses a compressed packet header and involves less overhead than SLIP. See also *SLIP*.

CSPDN—Circuit-Switched Public Data Network. A CSPDN is a communications network that uses circuit-switched digital data circuits and is available to the public. See also *circuit switching*, *digital signal*.

CSU—Channel Service Unit. Along with a Data Service Unit (DSU), a CSU is a component of Data Circuit-terminating Equipment (DCE). A CSU connects a digital phone line to a customer's network-access equipment. It can be built into the network interface of the network-access equipment, or it can be a separate device. The CSU terminates the connection at the user's end and processes digital signals. It also prevents a faulty DSU from interfering with data transmissions on the digital line. See also *DCE*, *digital signal*, *DSU*.

CTS—Clear To Send. CTS is a signal sent from a receiving device to a transmitting device, indicating that the transmitter can begin sending data. A CTS signal is generally a response to a transmitter's Request To Send (RTS) signal. See also *RTS*.

CUD—Call User Data. The CUD field identifies the encapsulation in use over an X.25 Virtual Circuit (VC). See also *encapsulation*, *VC*, *X.25*.

CUG—Closed User Group. A CUG is a calling group to which access is restricted. A user can be a member of more than one CUG. In general, members of a specific CUG can communicate among themselves, but not with users outside the group. In some cases, however, specific CUG members can originate calls to destinations outside the group, or receive calls from outside the group. The Network Service Provider (NSP) can determine the maximum number of CUGs a user can belong to. See also *CUG index*, *NSP*.

CUG index—Closed User Group index. On an X.25/PAD call, the CUG index indicates to the called switch the CUG selected for a virtual call. See also *CUG*, *X.25/PAD*.

Customer Premises Equipment—See *CPE*.

Cyclic Redundancy Check—See *CRC*.

D

D4-framed T1 line—A T1 line that uses the D4 format, also known as the *Superframe format*, to frame data at the physical layer. The D4 format consists of 12 consecutive frames, each one separated by framing bits. T1 lines that do not use ISDN D-channel signaling use the D4 format. See also *T1 line*.

D-A conversion—Digital-to-Analog conversion. D-A conversion is a process in which a digital signal is modified into an analog signal. D-A conversion takes place, for example, when digital data reaches an analog modem. Compare with *A-D conversion*. See also *analog signal*, *digital signal*, *modem.*

DASS-2—A signaling protocol used by British Telecom on ISDN links. DASS-2 specifies the signaling that occurs on the D channel. Although DASS-2 is widely available, many new installations use Q.931. See also *Q.931*.

Database-Description packet—A Type-2 Open Shortest Path First (OSPF) packet. OSPF routers exchange Database-Description packets when an adjacency is being initialized. Each packet describes the contents of the link-state database. The routers use a poll-response procedure. One of the routers is the master, and the other a slave. The master sends Database-Description poll packets, and the slave sends Database-Description response packets. OSPF links the responses to the polls by means of a sequence number in each packet. See also *adjacency*, *link-state database*, *OSPF*.

Data Carrier Detect—See *CD*.

Data Circuit-terminating Equipment—See *DCE*.

Data Communications Equipment—See *DCE*.

data compression—A method of reducing the size of a file in order to increase the data transmission rate. Compression algorithms, such as those incorporated in the V42bis standard, take advantage of redundancies in data files by substituting a few characters for many. Data compression is especially effective with text files and certain graphics file formats, and has become an important concern for network administrators. See also *compression*, *link compression, slot compression*, *Stac compression*, *Stac-9 compression*, *Stac LZS compression*, *V.42bis*, *VJ compression*.

Data Encryption Standard—See *DES*.

data filter—A packet filter that defines which packets the MAX can transmit on a connection. When you apply a data filter, its forward or drop action affects the actual data stream by preventing certain packets from reaching the Ethernet from the WAN, or vice versa (Figure 11).

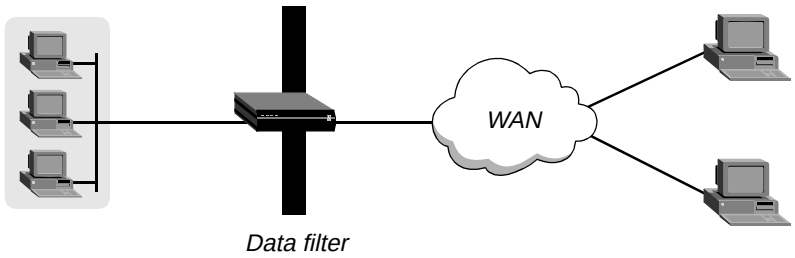

Figure 11. Data filters can drop or forward certain packets

Many sites use data filters for security purposes, but you can apply data filters to any purpose that requires the MAX to drop or forward only specific packets. For example, you can use data filters to drop packets addressed to particular hosts or to prevent broadcasts from going across the WAN. You can also use data filters to allow users to access only specific devices across the WAN. Compare with *call filter*. See also *generic filter, input filter, IP filter, IPX filter, output filter, packet filter.*

data frame—See *general frame.*

Datagram Delivery Protocol—See *DDP*.

datalink—The link interface to a Frame Relay device. The datalink refers to specific nailed-up bandwidth on the MAX and defines the operations and link-management functions that the MAX performs on the interface. See also *Frame Relay, Frame Relay network.*

Data Link Connection Indicator—See *DLCI.*

Data Link layer—The second layer of the OSI Reference Model. The Data Link layer creates, sends, and receives data packets appropriate for the type of network in use. Data-Link-layer protocols include High-Level Data Link Control (HDLC), Link Access Procedure, Balanced (LAPB), Link Access Procedure, D channel (LAPD), Point-to-Point Protocol (PPP), and Serial Line Internet Protocol (SLIP). See also *HDLC, LAPD, OSI Reference Model, PPP, SLIP*.

Data Over Subscriber Bearer Service—See *3.1 Khz audio-bearer service.*

data-over-voice—A method of sending digital data over telephone trunks by means of either voice-bearer service or 3.1 Khz audio-bearer service. See also *3.1 Khz audio-bearer service*.

Data Over Voice Bearer Service—See *3.1 Khz audio-bearer service*.

data rate—The transmission speed of data over a line, generally expressed as thousands of bits per second (Kbps).

data service—A service provided over a WAN line and characterized by the unit measure of its bandwidth. A data service can transmit either data or digitized voice. The following types of data services are available: Switched-56, Switched-64, Switched-384 (also known as *H0*), Switched-1536 (also known as *H11*), MultiRate, and GloBanD. See also *GloBanD*, *MultiRate*, *Switched-56*, *Switched-64*, *Switched-384*, *Switched-1536*.

Data Service Unit—See *DSU*.

Data Set Ready—See *DSR*.

Data Terminal Equipment—See *DTE*.

data-transfer mode—For an X.25/T3POS connection, the method used for error recovery and data transmission. The MAX enables you to specify Local, Transparent, Blind, and Binary Local. See also *Binary Local mode*, *Blind mode*, *Local mode*, *Transparent mode*.

Data Transmit Ready—See *DTR*.

DB25 pin connector—A 25-pin connector on which the RS-232C standard is based. Ten connections are commonly used. The names of the signals and pin designations on a standard DB25 pin connector are:

Pin	Signal
1	Protective (frame) ground
2	Transmit Data (TD)
3	Receive Data (RD)
4	Request To Send (RTS)
5	Clear To Send (CTS)
6	Data Set Ready (DSR)
7	Signal Ground (SG)
8	Carrier Detect (CD)
20	Data Terminal Ready (DTR)
22	Ring Indicator (RI)

See also *CD*, *CTS*, *DSR*, *DTR*, *protective ground*, *RD*, *RI*, *RS-232C*, *RTS*, *SG*, *TD*.

DBA—Dynamic Bandwidth Allocation. DBA denotes the process of adding or subtracting bandwidth from a switched connection in real time without terminating the link. Multilink Protocol Plus (MP+) and the Ascend Inverse Multiplexing (AIM) protocol support DBA based upon a set of parameters you specify. To add bandwidth, the MAX dials additional connections, and uses inverse multiplexing to add new channels to the call.

The MAX can reject a request to add bandwidth if no more channels are available, or if the network is congested. Under either of these conditions, the two ends enter bandwidth-addition-lockout mode, in which neither side can request bandwidth. The lockout prevents both ends from continually trying to add new channels unsuccessfully. The MAX and the Ascend unit at the other end of the link automatically remove the lockout restriction when the conditions that caused the lockout change. Changes typically result from plugging in a new switched-service line, reconfiguring a Line profile, or experiencing a switched-service congestion timeout. When the lockout ends, each end is free to add bandwidth.

If you use a circuit between two locations to capacity 24 hours per day, using a nailed-up line is more cost effective than using a switched line. However, if you need the circuit only sporadically, or if the circuit is sometimes underutilized, it often makes more sense to lease a smaller amount of nailed-up bandwidth and then supplement it with additional switched bandwidth as traffic requirements dictate.

For example, you might establish some connections only when you need to transfer data, and a single circuit can accommodate low traffic levels. However, if traffic levels grow beyond the capacity of the circuit (such as during a large file transfer), DBA automatically adds additional switched channels. When traffic levels subside, DBA automatically removes the channels from the connection. The bandwidth and connection costs are thereby reduced. You pay only for bandwidth when you need it.

See also *bandwidth*, *circuit*, *MP+*, *nailed-up line*, *switched line*.

DBM database—A RADIUS users file in UNIX database format. Compare with *flat ASCII users file.* See also *users file*.

DCD—See *CD*.

DCE—Data Circuit-terminating Equipment (also Data Communications Equipment). A DCE is a device that connects Data Terminal Equipment (DTE) to a communications channel, such as a telephone line. A DTE refers to a device that an operator uses, such as a computer or a terminal. A DCE converts the format of the data coming from the DTE into a signal suitable to the communications channel. An example of a DCE is a modem, which converts digital data from a computer to analog signals suitable for sending over a telephone line. See also *analog signal*, *digital data*, *DTE*, *modem*.

DCE interface—A MAX interface that provides AIM/BONDING inverse multiplexing services to a device connected to it.

D channel—A channel that carries WAN synchronization and signaling information on a T1 PRI or E1 PRI line. See also *E1 PRI line*, *T1 PRI line*.

DCS 1800—Digital Cellular System working at 1800 MHz. DCS 1800 is a European mobile-telephone service based on European Telecommunications Standards Institute (ETSI) standards. See also *ETSI*.

DCS 1900—See *GSM 1900*.

DDP—Datagram Delivery Protocol. DDP is an AppleTalk Network-layer protocol. It provides connectionless service between sockets, and handles both addressing and routing. See also *routing*, *socket*.

DE—Discard Eligibility. DE is a bit in a Frame Relay packet header. You set the DE bit to indicate that the network can discard the packet when traffic reaches a high level. See also *Frame Relay*.

default gateway—The default router the Ascend unit uses for traffic from a specific connection if it finds no explicit route in the IP routing table. See also *IP router*, *IP routing table*.

default route—The route the Ascend unit uses if it does not find a match for a packet's destination address. The default route has the destination address 0.0.0.0. If the Ascend unit finds a default route, it brings up the required connection (if necessary) and forwards the packet.

Figure 12 shows a router on a local subnet configured as the default route in a MAX. This type of configuration enables the MAX to turn off RIP on its local interfaces, and forward all local packets to the default route.

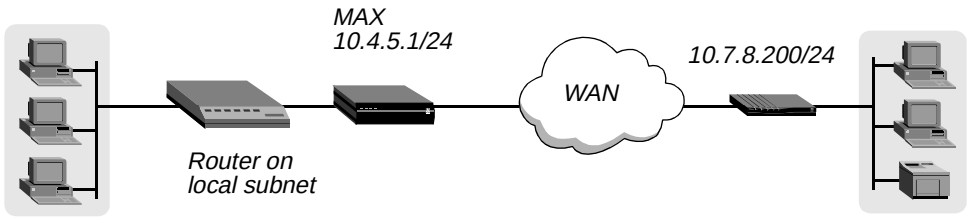

Figure 12. Default route to a local IP router

If the routing table has no default route and no route that matches a packet's destination address, the MAX drops the packet. See also *IP route*, *IP router*, *IP routing table*.

Default user name—A name the RADIUS server uses to grant access to clients who do not appear in the `users` file. You can configure only one Default profile. It must specify the user name Default, and it must be the last profile in the `users` file. See also *user profile*, *users file*.

default zone—The zone assigned to an AppleTalk service on an interface if the service does not select a zone in which to reside. See also *zone*, *zone list*.

Defender authentication—A form of token-card authentication that makes use of the AssureNet Pathways Defender authentication server. When you configure Defender authentication, the MAX forwards any call not authenticated by a local Connection profile to the Defender server. Defender authentication proceeds in three stages:

Stage	Description
1	Stage 1 begins a short time after the caller connects to the MAX, and before the MAX receives the first prompt from the authentication host. The Defender server provides the text of the prompts or challenges, and the MAX passes them to the caller. Calls in Stage 1 are preserved if an authentication host is unavailable or loses its connection. This situation might occur when the very first caller is authenticating with Defender after the router boots up, and the first authentication host is unavailable. The router tries the second and third hosts in order to authenticate the user.
2	Stage 2 occurs during the time the caller is interacting with the authentication host, but before the authentication sequence is complete. The Defender uses a challenge-response protocol, with a token card to provide the responses. Calls in Stage 2 are never preserved if an authentication hosts loses its connection. Defender has no mechanism for having one authentication server take over for another if the first loses a connection in the middle of a state.
3	Stage 3 occurs when the caller has completed authentication and is interacting with the MAX. Callers in Stage 3 are not dropped by the router, because their calls are already authenticated. However, if the host on which they authenticated is no longer available, their logout time is not sent (as would be the case if the host had remained connected). Defender provides no mechanism to notify one authentication host when a user call authenticated by another host is terminated.

You can use a Defender server with or without RADIUS authentication. The Defender server does not provide per-user control, such as enforcing a maximum number of channels. It provides only per-user authentication. If you need both per-user control and authentication, use RADIUS. See also *authentication server*, *RADIUS*, *token-card server*.

Denial of Service attack—See *DoS attack.*

density enforcement—For AMI-encoded T1 lines, requirements that dictate that you cannot transmit 16 consecutive zeroes. See also *AMI*, *T1 line*.

DES—Data Encryption Standard. DES is the U.S. encryption standard for nonclassified documents. This standard uses a 64-bit key and private-key encryption. In private-key encryption, only the sender and receiver know the key for encrypting the data. DES cannot ensure that the sender and receiver are legitimate. A sender who has learned the key can fraudulently use it. See also *encryption*, *private-key encryption.*

Designated Router—See *DR.*

DeskDial client—Ascend client software for network modem-pool access.

destination address—In a frame, packet, or message sent over a bridged or routed connection, the IP, IPX, AppleTalk, or hardware address of the intended recipient of the transmission. Compare with *source address*. See also *AppleTalk routing*, *hardware address*, *IP address*, *IP routing*, *IPX bridging*, *IPX routing*.

destination port—The port to use to communicate with the destination machine, such as a User Datagram Protocol (UDP) port on an authentication server, or an Simple Mail Transfer Protocol (SMTP) port on a mail server. Compare with *source port*. See also *SMTP*, *UDP*, *UDP port*.

Destination Service Access Point—See *DSAP*.

detail file—A file containing RADIUS accounting records. See also *accounting*.

device integration—In ISDN technology, the ability to carry digital signals from various devices over a single network interface. Communication that once required numerous wire pairs can now take place over a single wire pair by passing digital signals through local B channels. You can connect a variety of devices to an ISDN line by means of an ISDN Network Termination One (NT1) device or an ISDN Integrated Access Device (IAD). See also *IAD*, *ISDN*, *ISDN line*, *NT1*.

DHCP—Dynamic Host Configuration Protocol. DHCP is a TCP/IP protocol that enables a client to obtain a temporary IP address from a central server (known as a *DHCP server*). See also *DHCP server*, *dynamic IP*.

DHCP server—Dynamic Host Configuration Protocol server. A DHCP server assigns a temporary IP address to a client that requests it. See also *DHCP*, *DHCP spoofing*, *IP address*.

DHCP spoofing—Dynamic Host Configuration Protocol spoofing. A process that enables a local device to receive an IP address from a DHCP server across a slow WAN link. When you set up a DHCP connection, the MAX can assign a dynamic IP address to a remote DHCP client over a bridged connection. The MAX becomes a DHCP server.

For example, suppose a group of DHCP clients resides on a LAN connected to a Pipeline, and the Pipeline connects to the MAX over a bridged PPP connection (Figure 13). The MAX can assign dynamic IP addresses to any of the DHCP clients on the remote LAN.

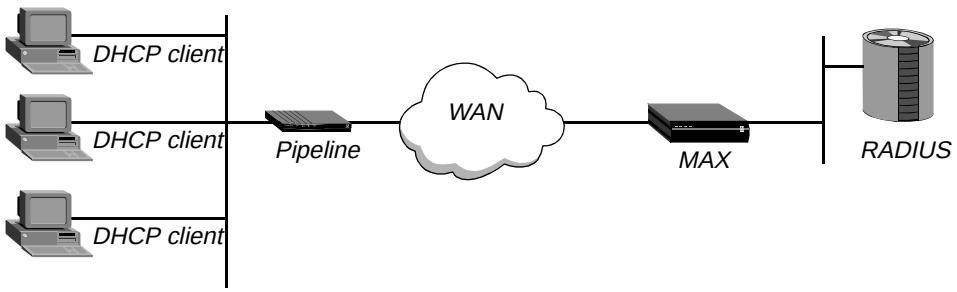

Figure 13. Pipeline connected to DHCP clients

The RADIUS server holds the configuration information the MAX uses to identify and authenticate each DHCP client. When a PC sends a broadcast DHCP request, the following events take place:

1. Acting as a DHCP server, the Pipeline receives the DHCP request, and sends the PC a temporary IP address. The address can be static or dynamic. It has a very short time-to-live (ttl).
2. The Pipeline dials the remote side, passing along the original DHCP request.
3. The DHCP server sends back a server-assigned IP address.
4. When the Pipeline receives the address from the remote side, it passes the address to the PC.
5. The PC changes its IP address to the server-assigned address.

Typically, a device requesting an IP address from a DHCP server waits a limited amount of time before timing out the request. For complex WAN links, with authentication processes, there may not be enough time to complete the process.

The DHCP server allocates an IP address from one of its IP address pools and assigns it to the client for 30 minutes. The client must renew the IP address assignment before the 30-minute period expires. The DHCP server uses its local memory to keep track of all IP addresses it has assigned. Therefore, it loses the entries for current, unexpired IP address assignments when you reset it.

A client can hold an unexpired IP address assignment when you reset the DHCP server. After the reset, the server might assign that address to a new client. The duplicate IP addresses cause network problems until the first assignment expires or one of the two clients reboots.

Compare with *IP address spoofing*, *IPX spoofing*, *SPX spoofing*, *watchdog spoofing*. See also *DHCP*, *DHCP server*, *IP address*.

diagnostic field—On an X.25 network, an optional field in a Clear-Request, Reset-Request, or Restart-Request packet. Each packet has a required Cause field and an optional Diagnostic field. If the Cause field indicates that the remote Data Terminal Equipment (DTE) did not request the clear, reset, or restart, the Diagnostic field has standard values. If the Cause field indicates that the remote DTE requested the clear, reset, or restart, the Diagnostic field contains information specified in the Cause field by the remote DTE. See also *cause field*, *Clear-Request packet*, *DTE*, *Reset-Request packet*, *Restart-Request packet*, *X.25*.

Dialed Number Information Service—See *DNIS*.

dial-in modem access—The ability of a remote client to use a modem to dial into a remote-access server at a Point Of Presence (POP) or central site. Compare with *dial-out modem access*.

dial-in user—A remote user or device that calls the MAX over a switched circuit and requests a connection.

dial-out modem access—The ability of a user at a workstation on a corporate LAN to dial out, by means of a shared modem, to a Point Of Presence (POP) or branch office. Compare with *dial-in modem access*.

dial plan—See *extended dial plan*.

dial query—A MAX feature designed for sites that support many clients and connections to only a few remote IPX networks. When it receives a SAP query for a file server (service type 0x04) and its SAP table has no entry for that service type, the MAX brings up all connections that enable dial query.

dial-up line—A connection or circuit between two sites through a switched telephone network. A dial-up line is most commonly associated with a voice telephone call between two locations. For modem access, a dial-up line forms a link between two distant computers or LANs. Not just restricted to landline connections, a dial-up connection can also be established through a circuit-switched cellular network. See also *cellular communication, circuit-switched line*.

dictionary file—The Ascend RADIUS dictionary. The `dictionary` file contains a list of all the attributes the RADIUS daemon supports, along with the possible values for each attribute. See also *RADIUS, RADIUS daemon, RADIUS server*.

digital cellular—See *PCS*.

digital data—Data that can have only a limited number of separate values. The time of day represented by a digital clock, or the temperature represented by a digital thermometer are examples of digital data. The digital values do not change continuously, but remain at one discrete value and then change to another discrete value. Compare with *analog data*. See also *digital signal*.

digital line—A line that transmits data by means of a digital signal. See also *digital signal*.

digital loopback—A procedure that tests the digital processing for a communications device. Compare with *analog loopback*. See also *local loopback, loopback, remote loopback*.

Digital Loop Carrier—See *DLC*.

digital modem—An internal device in the MAX that enables it to communicate over a digital line with a station connected to an analog line. Incoming modem calls and incoming digital calls come over the same digital line. The MAX can accept an incoming call from the network either as a pure digital stream, or as a digital stream encoded by Pulse Coded Modulation (PCM). A PCM-encoded digital stream contains a digitized version of the analog waveform sent by a device attached to a modem.

The MAX can also convert outgoing data into analog waveforms, convert these waveforms to a PCM-encoded digital stream, and send them to the network over a digital line. The network presents the data to the receiving modem in analog form over an analog line. The data looks exactly as it would appear if it had been sent by an analog-based modem.

See also *analog line, digital line, modem, Series56 Digital Modem module.*

Digital Private Network Signaling System—See *DPNSS*.

digital signal—A type of signal that uses a limited number of discrete values to encode data transmitted over a wire. The value of the data encoded in a digital signal depends upon the state of the signal during a particular time period. Therefore, the sender and the receiver must synchronize their clocks. Each clock runs at a baud rate, the number of times per second the state of the signal is read or set. Several clocking schemes are available, and digital signals often include clock timing cues. A digital signal can transmit analog or digital data. For example, a Compact Disc (CD) encodes music data into digital signals, while the wires between computers transmit digital data in digital signals. Compare with *analog signal*. See also *analog data, digital data*.

Digital Signal Processor—See *DSP*.

Digital Subscriber Line—See *DSL*.

digital-to-analog conversion—See *D-A conversion*.

direct-access dialout—A feature that enables terminal-server users to have direct access to a particular Telnet port for modem dialout. See also *modem dialout*.

direct route—A route that can reach a destination without going through any intervening routers. See also *route, router*.

Discard Eligibility—See *DE*.

Disconnect message—See *ISDN Disconnect message*.

Disconnect-Request packet—A message from a client of the MAX, asking the MAX to disconnect the session. See also *Disconnect-Request-ACKed packet*, *Disconnect-Request-NAKed packet*.

Disconnect-Request-ACKed packet— A message the MAX sends to a client if it found at least one session to disconnect. Compare with *Disconnect-Request-NAKed packet*. See also *Disconnect-Request packet*.

Disconnect-Request-NAKed packet—A message the MAX sends to a client if it could not find a session to disconnect. Compare with *Disconnect-Request-ACKed packet*. See also *Disconnect-Request packet*.

disk capture feature—A feature that allows your terminal emulator to capture to disk the ASCII characters it receives at its serial port. See also *serial port, terminal emulator*.

distance-vector metric—A metric that uses a hop count to select the shortest route to a destination network. Routing Information Protocol (RIP) always uses the lowest hop count, regardless of the speed or reliability of a link. Compare with *link-state metric*. See also *RIP*.

DIX connector—See *AUI*.

DLC—Digital Loop Carrier. DLC is equipment that concentrates analog local loop lines, digitizing and multiplexing calls for transmission to the Central Office (CO). See also *CO, local loop, multiplexing*.

DLCI—Data Link Connection Indicator. A DLCI is a number between 16 and 991 that the Frame Relay administrator assigns. It identifies the logical endpoints of a Virtual Circuit (VC). A Connection profile or RADIUS user profile specifies a DLCI for each connection. The Frame Relay switch uses the DLCI to route frames through the network. The DLCI can change as frames pass through multiple switches. See also *Connection profile*, *Frame Relay switch*, *user profile*, *VC*.

DLE, EOT command—A clear-request command signal from the Data Terminal Equipment (DTE) on a T3POS PAD connection. The X.25/T3POS PAD clears the call when it receives a DLE, EOT command. See also *DTE*, *X.25/T3POS*.

DLE, EOT timer— A timer that indicates the maximum idle time the PAD allows for a T3POS call. The DLE, EOT timer applies only to Transparent and Blind mode. It is disabled in both Local and Binary Local mode. The DLE, EOT timer is also called the *T5 timer*. See also *Binary Local mode*, *Blind mode*, *Local mode*, *Transparent mode*, *X.25/T3POS*.

DNIS—Dialed Number Information Service. DNIS is a telephone company service that provides information about the called number, such as the name and location of the target user or device.

DNS—Domain Name System. DNS is a TCP/IP service for centralized management of address resolution. Using DNS, you can specify a symbolic name instead of an IP address. A symbolic name consists of a user name and a domain name in the format username@domain_name. The user name corresponds to the host number in the IP address. The domain name corresponds to the network number in the IP address. A symbolic name might be steve@abc.com or joanne@xyz.edu. The domain identifier is the last part of the domain name, and identifies the type of organization to which the host belongs.

DNS maintains a database of network numbers and corresponding domain names. When you use a symbolic name, DNS translates the domain name into an IP address, and sends it over the network. When the Internet service provider receives the message, it uses its own database to look up the user name corresponding to the host number.

See also *DNS List Attempt*, *domain name*, *host number*, *IP address*, *IP network number*, *local DNS table*.

DNS List Attempt—A feature that enables the MAX to avoid tearing down physical links when a host is unavailable. Domain Name System (DNS) can return multiple addresses for a host name in response to a DNS query, but it does not include information about availability of the hosts. A user typically attempts to access the first address in the list. If that host is unavailable, the user must try the next host, and so forth. However, if the access attempt occurs automatically as part of immediate services, the physical connection is torn down when the initial connection fails. When you enable the DNS List Attempt feature, the MAX tries one entry in the DNS list of hosts, and if that connection fails, tries the next entry, and so on, without losing the WAN session. See also *DNS*.

domain identifier—The portion of a domain name that appears last and specifies the type of organization to which the host belongs. The Internet Network Information Center (InterNIC) provides the following domain identifiers:

Domain identifier	Description
.arpa	ARPANET
.com	Commercial enterprise
.edu	Educational institution
.gov	Governmental organization
.mil	Military organization
.org	An organization not covered by the other categories

domain name—The portion of a symbolic name that corresponds to the network number in the IP address. In the symbolic name steve@abc.com, the domain name is *abc.com*. See also *IP address, IP network number*.

Domain Name System—See *DNS*.

DO menu—A context-sensitive list of commands that appears when you press Ctrl-D. The commands in the DO menu vary depending upon the context. For example, if you press Ctrl-D in a Connection profile, the DO menu looks like this one:

```
DO…
>0=ESC
 1=Dial
 P=Password
 S=Save
 E=Termserv
 D=Diagnostics
```

The MAX supports the following DO commands:

Command	Description
Answer (DO 3)	Answer an incoming call.
Beg/End BERT (DO 7)	Begin/End a byte-error test.
Beg/End Rem LB (DO 6)	Begin/End a remote loopback.
Beg/End Rem Mgm (DO 8)	Begin/End remote management.
Close TELNET (DO C)	Close the current Telnet session.
Contract BW (DO 5)	Decrease bandwidth.
Diagnostics (DO D)	Access the diagnostic interface.
Dial (DO 1)	Dial the selected or current profile.
ESC (DO 0)	Abort and exit the DO menu.
Extend BW (DO 4)	Increase bandwidth.

Command	Description
Hang Up (DO 2)	Hang up from a call in progress.
Load (DO L)	Load parameter values into the current profile.
Menu Save (DO M) 8	Save the vt100 interface menu layout.
Resynchronize (DO R)	Resynchronize a call in progress.
Save (DO S)	Save parameter values into the specified profile.
Password (DO P) 9	Log into or out of the MAX.
Termserv (DO E)	Access the terminal server interface.
Toggle (DO T)	Toggle the Palmtop Controller.

DoS attack—Denial of Service attack. A DoS attack is a deliberate attempt to interfere with network performance by means of forged Internet Control Message Protocol (ICMP) Echo Request packets directed to IP broadcast addresses.

Under ordinary circumstances, in order to determine whether a machine on the Internet is connected and responding, a host sends an ICMP Echo Request packet. If a machine receives the packet, it returns an ICMP Echo Reply packet. In a DoS attack, however, an attacker directs ICMP Echo Request packets to IP broadcast addresses from one or more remote locations. An intermediary receives an ICMP Echo Request packet directed to the IP broadcast address of its network. If the intermediary does not filter ICMP traffic directed to IP broadcast addresses, the machines on the network receive the ICMP Echo Request packet and send an ICMP Echo Reply packet in return. The packets do not use the IP address of the source machine as the source address. Instead, they contain the spoofed source address of the intended victim. When all the machines at the intermediary's site respond to the ICMP Echo Requests, they send replies to the victim's device. An attacker can send DoS attacks to multiple intermediaries at the same time, causing all of the intermediaries to direct responses to the same victim.

Both the intermediary and victim of a DoS attack can suffer severely degraded network performance. To protect against DoS attacks, you should disable IP-directed broadcasts on the MAX. By disabling these broadcasts, you deny an attacker the ability to direct IP broadcast traffic onto your network. In addition, you should prevent the MAX from responding to ICMP packets sent to IP broadcast addresses. If someone compromises a machine on your network, he or she may try to launch an attack using the MAX as an intermediary, sending the ICMP Echo Request packet to the IP broadcast address of the local network. Because this traffic does not travel through a router to reach the machines on the local network, disabling IP-directed broadcasts on the MAX is not sufficient to prevent a DoS attack. You must also prevent the MAX from responding to ICMP packets sent to the local broadcast address. See also *Echo, ICMP*.

dotted decimal notation—A system for specifying an IP address or subnet mask. In dotted decimal notation, each of the four portions of the IP address or mask is separated from the others by a decimal point, as in the address 200.10.5.1. See also *IP address, subnet mask.*

DOVBS—See *3.1 Khz audio-bearer service.*

downstream path—The path a call takes from a carrier's Central Office (CO) to the end user's home.

DPNSS—Digital Private Network Signaling System. DPNSS is a standard that defines how different Private Branch Exchange (PBX) systems can interoperate to produce a single virtual PBX. See also *PBX*.

DR—Designated Router. The DR is the router with which all other Open Shortest Path First (OSPF) routers in a broadcast network establish adjacencies. Figure 14 illustrates a configuration with both a DR and a Backup Designated Router (BDR).

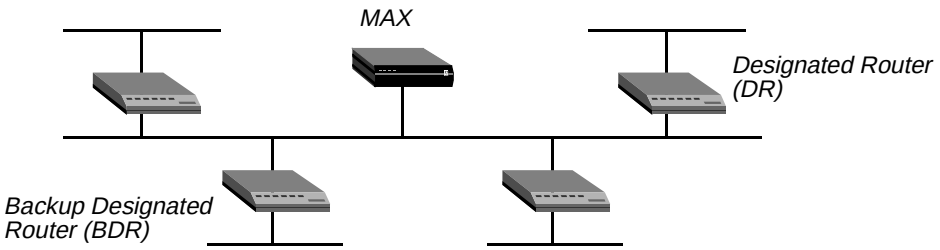

Figure 14. Designated and Backup Designated Routers

To reduce the number of adjacencies each router must form, OSPF calls one of the routers the DR. Doing so simplifies the routing table update procedure and reduces the number of link-state records in the database. The DR plays other important roles in reducing the overhead of OSPF link-state procedures. For example, other routers send Link-State Advertisements (LSAs) to the DR by using the "all-designated-routers" multicast address of 224.0.0.6.

The administrator chooses the DR based on the processing power, speed, and memory of the system, and then assigns priorities to other routers in case the BDR is down at the same time. The MAX can function as a DR or BDR. However, many sites choose to assign a LAN-based router as the DR or BDR in order to dedicate the MAX to WAN processing.

See also *adjacency*, *BDR*, *LSA*, *OSPF*, *router*.

DRAM—Dynamic Random Access Memory. DRAM is a kind of memory whose information resides in capacitors. The charge of each capacitor must be periodically refreshed. Compare with *EEPROM*, *NVRAM*, *RAM*.

DRAM card—On the MAX, a propriety Ascend card that provides DRAM to the system. The DRAM card is attached directly to the CPU bus, and should not be removed. Damage might occur if you attempt to remove it. See also *DRAM*, *DRAM interface.*

DRAM interface—On the MAX, an interface that accepts a plug-in DRAM card. See also *DRAM*, *DRAM card*.

Drop-and-Insert—A feature that enables a single T1 line to carry both data and voice traffic. The MAX allocates a preallocated portion of the T1 line to use both nailed-up and switched circuits for LAN internetworking. The remaining portion of the line can go to a PBX with a T1 interface. The PBX can access both nailed-up and switched circuits for voice purposes. You can also use Drop-and-Insert to share access-line bandwidth between the MAX and equipment other than a PBX, such as a channel bank or T1 multiplexer. See also *nailed-up circuit*, *PCM*, *switched circuit*, *T1 line.*

DS0 channel—A 64-Kbps D channel on a digital line.

DS0 minute—The online usage of a single 56-kbps or 64-kbps switched channel for one minute. For example, a 5-minute, 6-channel call uses 30 DS0 minutes.

DSAP—Destination Service Access Point. A DSAP is the Service Access Point (SAP) address at which the Logical Link Control (LLC) layer passes information to a Network-layer process. See also *SAP*, *SSAP*.

DSL—Digital Subscriber Line. DSL is a technology that provides high bandwidth over conventional copper wiring. See also *IDSL card*.

DSP—Digital Signal Processor. A DSP analyses and processes analog signals, converting them to a digital format. See also *analog signal*, *digital signal*.

DSR—Data Set Ready. DSR is a signal a modem transmits when it is ready to send and receive data.

DSU—Data Service Unit. Along with a Channel Service Unit (CSU), a DSU is a component of Data Circuit-terminating Equipment (DCE). The DSU connects to Data Terminal Equipment (DTE) by means of a synchronous serial interface, such as a V.35, RS-422, or RS-423 connection. The DSU formats and controls the flow of digital data between the network and the CSU. See also *CSU*, *DCE*, *digital data*, *DTE*, *RS-422*, *RS-423*, *V.35*.

DTE—Data Terminal Equipment. A DTE is a device that an operator uses, such as a computer or a terminal. Compare with *DCE*.

DTMF—Dual-Tone Multifrequency. DTMF is a technology enabling a touch-tone telephone to create 16 tones by means of 8 frequencies.

DTR—Data Transmit Ready. DTR is a signal indicating that a device is ready to transmit and receive data.

dual IP—A method of assigning a second IP address to the Ethernet interface in order to give the MAX a logical interface on two networks or subnets on the same backbone. See also *IP*, *IP address*.

dual-port call—A call in which a codec performs inverse multiplexing on two channels in order to achieve twice the bandwidth of a single channel. The codec provides two ports, one for each channel. Two AIM ports on the MAX connect a dual-port call to the codec. These ports are the primary port and the secondary port. Because the MAX places the two calls in tandem and clears the calls in tandem, it considers them a single call. See also *AIM port*, *codec*, *inverse multiplexing*.

Dual-Tone Multifrequency—See *DTMF*.

Dynamic Bandwidth Allocation—See *DBA*.

dynamic bandwidth overflow—A method of supplementing bandwidth during periods of peak demand. Through the mechanism of inverse multiplexing, the Ascend unit adds bandwidth when traffic reaches a preassigned level. See also *DBA*.

Dynamic Host Configuration Protocol—See *DHCP*.

dynamic IP—The process of assigning an IP address to a dial-in caller from an IP address pool. Figure 15 shows the MAX assigning an address from one of its defined pools to a dial-in host.

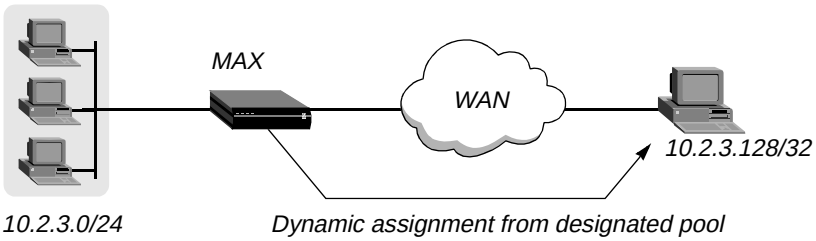

Figure 15. Dial-in host requiring assigned IP address

See also *IP address*, *IP address pool*.

Dynamic Random Access Memory—See *DRAM*.

dynamic route—A path to another network that the router learns by means of dynamic updates from other routers, rather than by means of a static specification in a configured profile. Routers that use Routing Information Protocol (RIP) broadcast their entire routing tables every 30 seconds, updating other routers about which routes are usable. Hosts that run Internet Control Message Protocol (ICMP) can also send ICMP Redirects to offer a better path to a destination network. Open Shortest Path First (OSPF) routers propagate link-state changes as they occur in order to update their routing tables. Compare with *multipath route*, *static route*. See also *IP route*, *IPX route*, *route*.

E

E1 line—A 2.048-Mbps line that supports 32 64-kbps channels, each of which can transmit and receive data or digitized voice. The line uses framing and signaling to achieve synchronous and reliable transmission. The most common configurations for E1 lines are E1 PRI and unchannelized E1. The MAX supports four E1 lines for up to 120 concurrent sessions. See also *E1 PRI line, unchannelized service*.

E1 PRI line—E1 Primary Rate Interface line. An E1 PRI line consists of 32 64-Kbps channels. It uses 30 B channels for user data, 1 64-Kbps D channel for ISDN D-channel signaling, and one framing channel. The B channels can be all switched, all nailed up, or a combination of switched and nailed up. The E1 PRI line is a standard in Europe and Asia called CEPT G.703. Compare with *ISDN BRI line, T1 PRI line, unchannelized service*. See also *B channel, D channel, E1 line, G.703, ISDN D-channel signaling, nailed-up channel, switched channel*.

E1 Primary Rate Interface line—See *E1 PRI line*.

E1-R2 Israeli signaling—See *R2 signaling*.

E2 line—An 8.45-Mbps line that supports four 2.048-Mbps E1 channels.

E3 line—A 34-Mbps line that supports 16 2.048-Mbps E1 channels.

EAS—External Authentication Server. See *authentication server*.

EAZ—Endgeraete Auswahl Ziffer. EAZ is a one-digit string appended to the phone number for calls over German 1TR6 lines.

Echo—A signal that determines whether a node can receive and acknowledge data transmissions. A host sends an Echo Request packet. If the destination is properly connected and receives the Echo Request packet, it sends back an Echo Reply packet.

echo cancellation—A method that the telephone company uses to remove echoes from an analog line. See also *analog line, echo tail*.

echo tail—The amount of hybrid-line echo that the G.165 echo canceller can eliminate in VoIP calls. The echo canceller is design to eliminate the echo generated when a voice call is transmitted across a two-wire/four-wire boundary. The echo canceller is applied to the speech on the trunk side of the MAX. It does not suppress the acoustic echo that is normally generated at the calling endpoint (receiver/transmitter echo). When the length of the hybrid-line echo exceeds 32ms, echo will be detected at the distant endpoint.

ECN—Explicit Congestion Notification. ECN is a method of informing Frame Relay nodes that there is traffic congestion on the network. The Frame Relay header can use a Backward Explicit Congestion Notification (BECN) bit or a Forward Explicit Congestion Notification (FECN) bit to notify nodes of traffic congestion. *BECN*, *FECN*, *Frame Relay*.

EEPROM—Electronically Erasable Programmable Read-Only Memory. EEPROM is a type of Programmable Read-Only Memory (PROM) that can be erased by exposing it to an electrical charge. It retains its contents across resets and power cycles, and is similar to NVRAM. With EEPROM, data is written or erased one byte at a time. With NVRAM, data is written or erased in blocks. See also *NVRAM*, *PROM*.

EGP—Exterior Gateway Protocol. EGP is a type of protocol used to exchange routing information between one Open Shortest Path First (OSPF) Autonomous System (AS) and another. The AS number may be used by Area Border Routers (ABRs) to filter out certain EGP routing information. OSPF can make use of EGP data generated by other border routers and added to the OSPF system as Autonomous System Externals (ASEs). See also *ABR*, *AS*, *ASE*, *OSPF*.

EIA—Electronic Industries Association. The EIA is a group that determines standards for electrical transmission.

EIA/TIA-232—A Physical-layer standard nearly identical to V.24. EIA/TIA-232 is also known as *RS-232*. See also *RS-232*.

EIA-449—A Physical-layer standard also known as *RS-449*. See also *RS-449*.

EIA-530—A way of referring to RS-422 and RS-423. See also *RS-422*, *RS-423*.

Electronically Erasable Programmable Read-Only Memory—
See *EEPROM*.

Electronic Industries Association—See *EIA*.

Embedded Operations Channel—See *EOC*.

en-bloc dialing—See *Q.931 en-bloc dialing*.

encapsulation—A technique used by layered protocols in which a low-level protocol accepts a message from a higher-level protocol, and then places the message in the data portion of the lower-level frame. The logistics of encapsulation require that packets traveling over a physical network contain a sequence of headers. Encapsulation enables the transmission of data over networks that use differing protocols.

encryption—A process that takes ordinary data and converts it into a format unreadable to anyone without a decryption key. Authorized personnel with access to this key can unscramble the information. Data encryption is a useful tool against network snoopers. See also *private-key encryption*, *public-key encryption*.

endpoint—When a tunneling protocol is in use, the system that encapsulates the packets (the Foreign Agent) or the system that decapsulates the packets (the Tunnel Server). Examples of tunneling protocols are Ascend Tunnel Management Protocol (ATMP), Layer 2 Tunneling Protocol (L2TP), and Point-to-Point Tunneling Protocol (PPTP). See also *ATMP*, *L2TP*, *PPTP*.

endpoint discriminator option—On an Multilink Protocol (MP) or Multilink Protocol Plus (MP+) link, a method of identifying the system transmitting the packet. The endpoint discriminator option tells a system that the peer on the link could be the same as the peer on another connection. If the option distinguishes the peer from all others, the system must establish a new bundle for the link. See also *bundle*, *MP*, *MP+*.

Enhanced Through Cellular—See *ETC*.

Enigma—An important provider of network security applications. Enigma's SafeWord AS (also known as the *Enigma Logic SafeWord server*) is a UNIX-based software authentication server that identifies users by means of dynamic passwords (called *tokens*). The server identifies users at the point of connection to a TCP/IP network, and uses standard network authentication protocols and token cards. See also *SafeWord authentication*, *SafeWord token*, *token card*.

Enigma Logic SafeWord server—A token-card authentication server from Enigma. See also *Enigma, SafeWord authentication, SafeWord token, token card, token-card authentication*, *token-card server*.

ENQ—A control character that signifies a request for identification or status on an X.25/ T3POS connection. See also *ENQ handling timer*, *X.25/T3POS*.

ENQ handling timer—A timer that indicates the amount of time the Packet Assembler/ Disassembler (PAD) waits for an ENQ from the host on an X.25/T3POS connection. See also *ENQ*, *PAD*, *X.25/T3POS*.

enterprise-wide network—A network that contains all or most of a company's hardware and software resources. Typically, an enterprise-wide network includes computers that run different operating systems and reside on different types of networks. Therefore, achieving interoperability is the biggest challenge facing the administrator of an enterprise-wide network.

EOC—Embedded Operations Channel. In the BRI-U interface, EOC is the out-of-band mechanism for implementing maintenance functions. Instead of using the D or B channels, EOC uses the maintenance bits of the U-interface superframe. Maintenance functions include test loopback, statistics gathering, and requests to generate errors (to check that the block-error counters work). You can perform EOC loopback on a B channel during a session with an IDSL card. See also *B channel*, *D channel*, *IDSL card*.

error correction—A process that determines whether line noise has caused data to be garbled or dropped in transit, and works to correct the problem. The two most common error-correction protocols and standards used by analog modems are MNP and V.42. See also *MNP*, *V.42*.

ESF—Extended SuperFrame. ESF is a framing format that consists of 24 consecutive frames, separated by framing bits. The ISDN specification advises that you use ESF with ISDN D-channel signaling. See also *ISDN D-channel signaling*.

ETC—Enhanced Throughput Cellular. An error-correction protocol developed by AT&T Paradyne, Inc. ETC is based on V.32bis, providing a maximum rate of 14,400 bps. See also *V.32bis*.

Ethernet—The most commonly used architecture for Local Area Networks (LANs), connecting devices such as computers, printers, and terminals. An Ethernet network uses the Physical and Data Link layers for data transmission. Ethernet incorporates a bus topology, and can operate at a rate of up to 10 Mbps. See also *Data Link layer*, *Physical layer*.

Ethernet II—A protocol specification for the Media Access Control (MAC) header of an IPX frame. Compare with *802.2*, *802.3*, *SNAP*. See also *IPX frame*, *MAC*.

Ethernet address—See *MAC address*.

Ethernet Status window—A status window that displays statistics about each active Ethernet interface. Following is a sample Ethernet status window:

```
|--------------------|
|90-400 Ether Stat   |
|>Rx Pkt:    3486092 |
| Tx Pkt:      10056 |
|    Col:       3530 |
|--------------------|
```

The Ethernet Status window shows the current count of received frames, transmitted frames, and frames with errors at the Ethernet interface. See also *status window*.

Ethernet transceiver—A device that connects workstations to standard thick or thin Ethernet-style cable. An Ethernet transceiver sends and receives information, and offers data-packet collision detection. See also *Thick Ethernet*, *Thin Ethernet*.

ETSI—European Telecommunications Standards Institute. ETSI is a European organization established in 1988 to provide common telecommunications standards.

EU-RAW—A WAN encapsulation protocol used primarily in Europe. IP packets are HDLC-encapsulated and include a Cyclic Redundancy Check (CRC).

European Telecommunications Standards Institute—See *ETSI*.

EU-UI—A WAN encapsulation protocol used primarily in Europe. IP packets are HDLC-encapsulated, and include a special header and a Cyclic Redundancy Check (CRC).

even parity—See *parity*.

exclusive port routing—A feature that causes the MAX to drop calls for which it has no explicit call-routing information (such as answer numbers, ISDN subaddressing, and so on). When exclusive port routing is disabled (the default), and the bearer service is voice, the MAX routes the call to a digital modem; if the bearer service is V.110, the MAX routes the call to the first available V.110 module; if the bearer service is data, the MAX routes the call to the first available AIM port; if no AIM port is available, the MAX routes the call to the bridge/router. See also *answer number*, *call routing*, *subaddress*.

expansion card—See *slot card*.

expansion module—See *slot card.*

expansion slot—See *slot*.

expect-send script—A script whose lines begin with either the `send` or the `expect` command. A line that begins with `send` causes all the other characters on the line to go through the WAN port running the script. A line that begins with `expect` causes the router to wait for matching characters from the WAN port. You can use an expect-send script to authenticate logins to the terminal server, or to start a Point-to-Point Protocol (PPP) or Serial Line Internet Protocol (SLIP) session from within a terminal-server connection. See also *authentication*, *PPP*, *SLIP*, *terminal mode*.

Explicit Congestion Notification—See *ECN*.

extended AppleTalk network—An AppleTalk Phase 2 network that uses an extended addressing scheme. With extended addressing, AppleTalk uses an 8-bit node number and a 16-bit network number for each host, allowing up to 16 million hosts on one network. See also *AppleTalk*, *AppleTalk routing*.

extended dial plan—Specifies how the device terminating an IDSL line can send and receive calls. The dial plan uses a specified trunk group, but accesses a Dial Plan profile to obtain T1 PRI parameters for the outgoing call. The extended dial plan is typically used to route calls from a terminating device on a Host BRI line out to the WAN using T1 PRI channels. However, it can also be used to set up the T1 PRI parameters for other outgoing calls. See also *IDSL card*, *T1 PRI line*.

extended load—A MAX release denoted by an `f` preceding the build name. For example, `fti.m40` is the extended load for the MAX 4000 T1 IP-only software build. You must use Trivial File Transfer Protocol (TFTP) to upgrade to extended loads. See also *TFTP*.

Extended SuperFrame—See *ESF*.

Exterior Gateway Protocol—See *EGP*.

external authentication—A remote method of identifying the users permitted to access network resources. The remote server can be a RADIUS, TACACS, TACACS+, or token-card server. See also *RADIUS*, *RADIUS server*, *TACACS*, *TACACS+*, *token-card server*.

External Authentication Server—See *authentication server*.

external LSA—External Link-State Advertisement. In an Open Shortest Path First (OSPF) configuration, an external LSA is exchanged between Autonomous Systems (AS) by Autonomous System Border Routers (ASBRs). Compare with *internal LSA*. See also *AS*, *ASBR*, *LSA*, *OSPF*.

external route—A route imported into the Open Shortest Path First (OSPF) database from outside the router's Autonomous System (AS). Compare with *intra-area route*. See also *AS*, *OSPF*, *route*.

F

facility—An optional service offered by an X.25 packet-switching network. The user can request a facility when subscribing for network service or when establishing a call. See also *X.25*.

Facilities Data Link—See *FDL*.

Failure-to-start record—A RADIUS-accounting or call-logging record that contains information about a failed login attempt. See also *Failure-to-start session*.

Failure-to-start session—An event denoting that a login attempt has failed. Information about this event appears in a RADIUS-accounting or call-logging Failure-to-start record.

fallback/fall-forward line sensing—A feature that enables a high-speed analog modem to monitor the quality of the phone line and step down to the next-lower speed if the line quality deteriorates. The modem falls forward to the next higher speed as line quality improves. See also *modem*.

Far-End Block Error—See *FEBE*.

Fast Ethernet—A LAN transmission standard with a data rate of 100 Mbps. Workstations with a 10-Mbps (10Base-T) Ethernet cards can be connected to a Fast Ethernet network.

fatal error—An error that causes the abrupt termination of a program.

fat load—Version 4.6Ci19 to 5.0Aix and all 5.0Ap releases with a file size greater than 960K. Before upgrading to a fat load for the first time, you must upgrade to a thin load. In addition, you must use Trivial File Transfer Protocol (TFTP) to upgrade to a fat load. Compare with *thin load*.

FDDI—Fiber Distributed Data Interface. FDDI is a proposed ANSI standard for a network architecture that uses high-speed fiber-optic lines and supports transmission rates of up to 100 Mbps.

FDL—Facilities Data Link. An FDL is a 4-Kbps digital link between a sender and the telephone company's monitors. The link uses Extended Superframe (ESF) framing. The telephone company uses an FDL to check on the quality and performance of T1 lines. See also *ESF*, *T1 line*.

FEBE—Far-End Block Error. FEBE is a signal the remote end sends to indicate that it has received E1 frames with either Framing Errors (FERR) or C-bit Parity Errors (CPERR). A block error is detected each time the calculated checksum of the received data does not correspond to the control checksum transmitted in the successive superframe. One block error indicates that one superframe has not been transmitted correctly. No conclusion with respect to the number of bit errors can be drawn from the block-error counter. Compare with *NEBE*. See also *CPERR*, *FERR*.

FECN—Forward Explicit Congestion Notification. FECN is a bit set in a Frame Relay header to notify a destination node that there is traffic congestion on the network. Compare with *BECN*. See also *Frame Relay*.

FERR—Framing Errors. FERR indicates the number of errors in the bits used to frame the E1 signal. See also *E1 line*.

Fiber Distributed Data Interface—See *FDDI*.

FIFO—First-In, First-Out. An algorithm that specifies that the first data in a buffer is the first data to be removed from the buffer.

File Transfer Protocol—See *FTP*.

filter—A set of rules describing what action the MAX should take when it encounters certain types of packets. A filter can apply to incoming packets, outgoing packets, or both. See also *packet filter*.

filter persistence—A method of enabling a firewall to persist across connection-state changes. With filter persistence, the firewall rules stay in force even when a connection goes offline. Filter persistence applies only to interfaces you configure on the MAX, such as Ethernet interfaces and virtual interfaces associated with Connection profiles. Filter persistence does not apply to interfaces built from RADIUS, TACACS, or TACACS+ profiles.

Filter profile—A profile containing parameters that set up filter rules. See also *filter*, *packet filter*.

Finger—A simple protocol that provides access to a Remote User Information Program (RUIP). Using the Finger protocol, the Finger utility can determine whether a particular user is logged into a certain device, and can gather other information about the user.

firewall—See *Secure Access Firewall*.

First-In, First-Out—See *FIFO*.

flag pattern—A pattern of 1s and 0s (01111110) that the MAX can use as the idle indicator on a dynamic AIM call. Compare with *mark pattern*.

flash card—See *PCMCIA flash card*.

flash memory—See *NVRAM*.

flat ASCII users file—A RADIUS users file in a flat, non-database format. Compare with *DBM database.* See also *RADIUS, users file*.

flow control—A method of compensating for differences in the flow of incoming and outgoing data for a modem or other device. If one system receives more data than it can hold in its buffers or process at any given time, it signals the sender to pause the transmission. Flow control can take place by means of hardware or software. See also *hardware flow control*, *software flow control*.

Foreign Agent—When a tunneling protocol is in use, the system that encapsulates the packets. Examples of tunneling protocols are Ascend Tunnel Management Protocol (ATMP), Layer 2 Tunneling Protocol (L2TP), and Point-to-Point Tunneling Protocol (PPTP).

When ATMP is in use, the Foreign Agent is an Ascend unit that a Mobile Client dials into. It is the starting point of the ATMP tunnel. Under ATMP, the Foreign Agent must be able to bring up an IP connection to the Home Agent, and it must authenticate the Mobile Client by means of a RADIUS user profile or Connection profile.

See also *ATMP*, *Home Agent, L2TP*, *PPTP*, *RADIUS, user profile*.

Forward Explicit Congestion Notification—See *FECN*.

Fractional T1—See *FT1*.

fractional T1 line—A nailed-up T1 line with bandwidth that might be only a fraction of the full T1 bandwidth. See also *T1 line*.

Fractional T1 Plus Switched Multilink Protocol Plus—See *FT1-MP+*.

frame—In Token Ring and Systems Network Architecture (SNA), a packet at the Data Link layer of the OSI Reference Model; in Frame Relay, a packet of fixed size; in Time Division Multiplexing (TDM), a sequence of time slots, each containing a portion of a multiplexed channel. A frame contains source and destination information, flags that designate the start and end of the frame, and information about the integrity of the frame. All other data, such as network protocol information and the actual payload of data, is first encapsulated in a packet. The system then encapsulates the packet in a frame. See also *Data Link layer*, *Frame Relay*, *OSI Reference Model*, *packet*, *TDM*.

framed protocol—A synchronous protocol that encapsulates data into frames. See also *framing*, *protocol*, *synchronous transmission*.

Frame Relay—A WAN architecture originally developed for ISDN lines. A Frame Relay network provides high throughput by handing monitoring functions to higher-level protocols. It is a very efficient standard, with a bandwidth of up to 2 Mbps. Frame Relay is ideal for situations in which periods of very high traffic are interspersed with idle periods. It is protocol independent, and performs routing over Virtual Circuits (VCs) called Data Link Connection Indicators (DLCIs). See also *DLCI*, *Frame Relay concentrator*, *Frame Relay Direct*, *Frame Relay network*, *Frame Relay profile*, *Frame Relay switch*, *Frame Relay user profile*, *ISDN*.

Frame Relay Annex A—The ITU standard defining Frame Relay maintenance procedures for Permanent Virtual Circuits (PVCs). ITU Annex A is analogous to the ANSI standard Annex D. See also *Frame Relay Annex D*.

Frame Relay Annex D—Part of the ANSI T1.617 standard. Annex D defines maintenance procedures that apply to Permanent Virtual Circuits (PVCs). Before ANSI standard T1.617 Annex D, a group of companies defined the Link Management Interface (LMI) mechanism for PVC management. The LMI specifies functionality similar to that defined by the ANSI standard, and is still widely supported. ANSI Annex D is analogous to the ITU standard Annex A. See also *Frame Relay*, *Frame Relay Annex A*, *LMI*, *PVC*.

Frame Relay concentrator—A device that concentrates many low-speed, dial-in connections into one high-speed, nailed-up connection to a Frame Relay switch. As a Frame Relay concentrator, the MAX forwards many lower-speed PPP connections onto one or more high-speed Frame Relay interfaces, as shown in Figure 16.

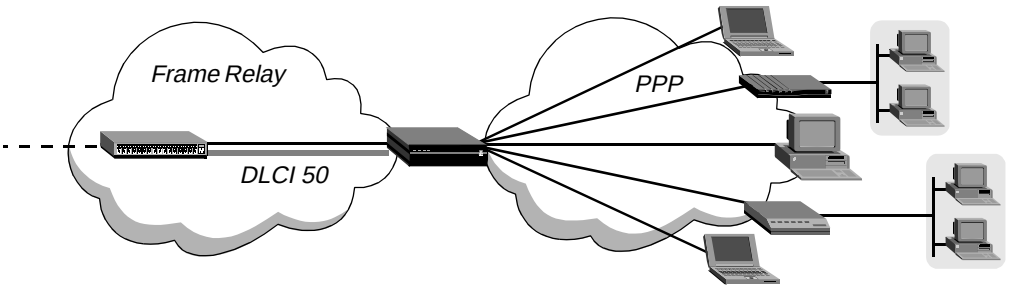

Figure 16. Frame Relay concentrator

In this kind of configuration, the decision to forward frames onto the Frame Relay interface can be made at OSI layer 3 (IP routing), or by Frame Relay Direct. Compare with *Frame Relay switch*. See also *Frame Relay*, *Frame Relay Direct*, *Frame Relay network*, *IP routing*.

Frame Relay Direct—A Frame Relay connection in which the MAX ignores the destination IP address in a packet from a dial-in PPP client, and uses the Data Link Connection Indicator (DLCI) to route the packet instead. In effect, the MAX does not route packets from the client in the usual sense. It simply passes them on to the Frame Relay network, and assumes that another device will route the packets on the basis of the destination IP address. Figure 17 shows two incoming PPP connections redirected out to the Frame Relay network. Both direct connections (shown at the right of Figure 17) use the same DLCI number (72).

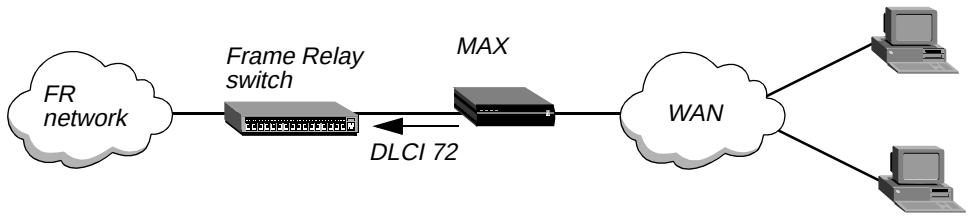

Figure 17. Frame Relay Direct connections using the same DLCI

A Frame Relay Direct connection is not a full-duplex tunnel between the PPP dial-in device and the switch. The MAX router handles the IP packets coming back from the Frame Relay switch, so the packets must contain the PPP caller's IP address for proper routing back across the WAN. See also *DLCI*, *Frame Relay*, *Frame Relay network*, *Frame Relay switch*, *IP address*, *PPP*.

Frame Relay gateway—A routing link in a Frame Relay configuration. A common way to concentrate incoming connections onto a Frame Relay link is to make use of OSI layer 3 (IP routing). For this purpose, the MAX requires ordinary profiles for callers, and a Data Link Connection Indicator (DLCI) logical interface that specifies a destination IP router. When a client dials in to reach the destination router, the MAX acts as a Frame Relay gateway, consulting its routing table to forward the packets onto Frame Relay.

In Figure 18, the MAX acts as a gateway between an IP router that dials in using PPP, and a remote router across the Frame Relay network.

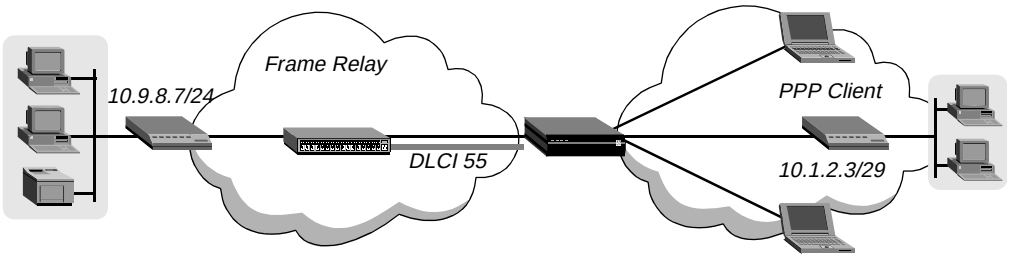

Figure 18. Frame Relay gateway

See also *DLCI*, *Frame Relay network*, *IP router*, *IP routing*, *IP routing table*.

Frame Relay link management—A feature that enables you to retrieve information about the status of the Frame Relay interface by means of special management frames with a unique Data Link Connection Indicator (DLCI) address. (DLCI 0 is the default for link-management frames.) On a User-to-Network interface (UNI) to Frame Relay, link-management procedures occur in one direction. The UNI-DTE device requests information, and the UNI-DTE device provides it. On a Network-to-Network interface (NNI), link-management procedures are bidirectional. Switches perform both the DTE and DCE link-management functions, because both sides of the connection request information from their peer switches. See also *DLCI*, *Frame Relay network*, *NNI*, *UNI*, *UNI-DCE interface*, *UNI-DTE interface*.

Frame Relay network—A network in which every access point connects directly to a Frame Relay switch. Depending on how a device such as the MAX is integrated into the Frame Relay network, it can operate as a Frame Relay terminating unit (Customer Premises Equipment or CPE) or as a Frame Relay switch.

A CPE is the source or destination of data traversing the Frame Relay service. For example, the MAX labeled MAX-02 in Figure 19 is the source and destination of the data stream from its PPP callers. When it is configured with a User-to-Network interface (UNI) to Frame Relay, the MAX acts as the user side (UNI-DTE) communicating with the network side (UNI-DCE) of a switch.

A network-side device connects the CPE device to a Frame Relay network. For example, the unit labeled MAX-01 in Figure 19 receives Frame Relay encapsulated frames from a CPE and forwards them on to another Frame Relay switch. When it is configured with a UNI-DCE interface, the MAX acts as the network side (UNI-DCE) communicating with the user side (UNI-DTE) of a Frame Relay device.

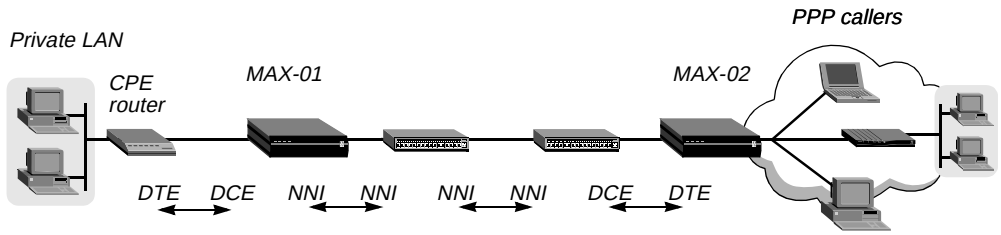

Figure 19. Frame Relay network

A Frame Relay switch is another kind of network-side device. It switches frames from one interface to another and exchanges status information with its peer switch. For example, the unit labeled MAX-01 in Figure 19 receives frames from its peer switch and switches them to its other Frame Relay interface. When it is configured with a Network-to-Network interface (NNI) to Frame Relay, the MAX acts as a Frame Relay switch. Switch-to-switch communication includes both user side (NNI-DTE) and network side (NNI-DCE) functions.

See also *CPE*, *Frame Relay*, *Frame Relay concentrator*, *Frame Relay Direct*, *Frame Relay profile*, *Frame Relay switch*, *Frame Relay user profile*, *NNI*, *UNI*, *UNI-DCE interface*, *UNI-DTE interface*.

Frame Relay profile—A profile that defines datalink operations, including link-management functions, for Frame Relay links. See also *datalink*, *Frame Relay*, *Frame Relay network.*

Frame Relay switch—A device that sends Frame Relay data out to the Frame Relay network. As a Frame Relay switch, the MAX receives frames on one interface and transmits them on another interface. The decision to forward frames onto the Frame Relay interface is made at OSI layer 2. The MAX router software is not involved.

To use the MAX as a switch, you must configure a circuit that pairs two Frame Relay interfaces. Instead of going to the layer 3 router for a decision on which interface to forward the frames, it relies on the circuit configuration to relay the frames received on one interface to its paired interface. A circuit is defined in two Connection or RADIUS profiles.

Figure 20 shows the MAX operating as a Frame Relay switch.

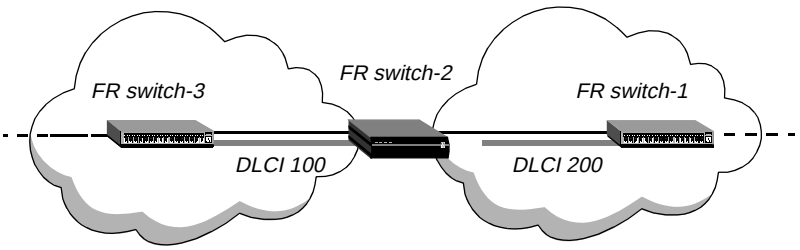

Figure 20. Frame Relay switch

Compare with *Frame Relay concentrator*. See also *Frame Relay*, *Frame Relay network.*

Frame Relay user profile—A RADIUS user profile that enables you to configure a Frame Relay connection for a user accessing a Frame Relay link. See also *Frame Relay*, *Frame Relay network*, *user profile*.

Frame window size—See *Level 2 Window Size*.

framing—At the Physical and Data Link layers of the OSI model, a method of fitting bits into a unit called a *frame.* A frame contains source and destination information, flags that designate the start and end of the frame, and information about the integrity of the frame. All other data, such as network protocol information and the actual payload of data, is first encapsulated in a packet. The system then encapsulates the packet in a frame. See also *Data Link layer*, *encapsulation*, *OSI Reference Model*, *packet*, *Physical layer.*

Framing Errors—See *FERR*.

FT1—Fractional T1. FT1 is a type of call that consists entirely of nailed-up channels. The call connects to a Terminal Adapter (TA), Channel Service Unit (CSU), or Data Service Unit (DSU) over fractional T1 or other nailed-up circuits. See also *CSU, DSU, fractional T1 line, FT1-AIM, FT1-B&O, FT1-MP+, nailed-up line, TA.*

FT1-AIM—Fractional T1-Ascend Inverse Multiplexing. FT1-AIM is a type of call in which the MAX combines nailed-up channels with switched channels to achieve the required bandwidth. An FT1-AIM call uses the AIM protocol, and is not available on host ports not equipped with AIM functionality. Both ends of the call must have AIM-compatible equipment. When the quality of a nailed-up channel falls to Marginal or Poor in an FT1-AIM call, the MAX drops the channel and does not replace it. The MAX cannot monitor these channels or restore them to an online call. See also *FT1, FT1-B&O, FT1-MP+*.

FT1-B&O—Fractional T1-Backup and Overflow. FT1-B&O is a type of call that provides automatic protection of nailed-up circuits.

- In providing backup bandwidth, the MAX drops all the nailed-up channels when the quality of a nailed-up channel falls to Marginal or Poor. The MAX then attempts to replace dropped nailed-up channels with switched channels. The MAX also monitors dropped nailed-up channels. When the quality of all dropped channels changes to Fair or Good, the MAX reinstates them.
- In providing overflow protection, the MAX supplies supplemental dial-up bandwidth during times of peak demand in order to prevent saturation of a nailed-up line. The circuit remains in place until the traffic subsides, and then it is removed.

The backup and overflow feature uses the Ascend Inverse Multiplexing (AIM) protocol, and is not available on host ports not equipped with AIM functionality. Both ends of the call must have AIM-compatible equipment. You must limit calls of this type to 28 channels. See also *FT1*, *FT1-AIM*, *FT1-MP+*.

FT1-MP+—Fractional T1 Plus Switched Multilink Protocol Plus. An FT1-MP+ connection begins as a nailed-up connection, but can later use switched channels, either to increase bandwidth or to provide a backup if the nailed-up channels go offline. When a nailed-up connection is temporarily down, the MAX polls continuously while trying to re-establish the link. If an outgoing packet arrives while the nailed-up connection is still down, the unit replaces the nailed-up channel with a switched channel. See also *fractional T1 line, FT1, FT1-AIM, FT1-B&O, MP+*.

FTP—File Transfer Protocol. FTP is an Application-layer protocol that enables you to transfer files from one device to another over a network. See also *Application layer*.

FTP server—A server that a user can contact in order to transfer files by means of the File Transfer Protocol (FTP) over a TCP/IP network.

full duplex—A type of communications configuration in which data can be transmitted in both directions at the same time. Compare with *half duplex*.

G

G.703—A standard specifying the physical and electrical characteristics of digital devices, including those at 64 Mbps and 2.048 Mbps.

G.711 A-Law—See *A-Law*.

G.711 U-Law—See *U-Law*.

gateway—A device or program that provides mapping at all seven layers of the OSI model and translates between two otherwise incompatible networks or network segments. A gateway performs code and protocol conversion to facilitate traffic between data highways of differing architectures. See also *OSI Reference Model*.

Gateway Home Agent—In an Ascend Tunnel Management Protocol (ATMP) configuration, a Home Agent that tunnels packets from the Foreign Agent to the Home Network across an open WAN connection. The WAN connection must be online. The Gateway Home Agent does not bring up a WAN connection to the Home Network in response to a packet it receives through the tunnel. For this reason, the Gateway Home Agent must have a nailed-up WAN connection to the Home Network. Compare with *Router Home Agent*. See also *ATMP, Foreign Agent, Home Agent, Home Network, nailed-up circuit*.

general frame—On an X.25/T3POS network, a frame defined as any sequence of octets received from or sent to the Data Terminal Equipment (DTE) within the period specified by the Char-to-Char timer. (A general frame is also known as a *data frame*.) In Local and Binary Local modes, and in opening frames, general frames are encapsulated in the format:

`<STX [`*`data`*`] ETX XRC>`

where:

- `STX` is the ASCII character \002.
- *`data`* is the user data being sent in this frame.
- `ETX` is the ASCII character \003
- `XRC` is the checksum. For all modes except Binary Local, the checksum is a one-character Longitudinal Redundancy Check (LRC) checksum. For Binary Local mode, the checksum is a two character Cyclic Redundancy Check (CRC) checksum.

Compare with *general frame, I-frame*. See also *Binary Local mode, Blind mode, Char-to-Char timer, CRC, DTE, LRC, X.25/T3POS*.

generic filter—A packet filter that examines the byte- or bit-level contents of a packet and compares them with a value defined in the filter. To use a generic filter effectively, you need to know the contents of certain bytes in the packets you wish to filter. Protocol specifications are usually the best source of such information. Compare with *IP filter, IPX filter*. See also *call filter, data filter, packet filter.*

Generic Routing Encapsulation—See *GRE*.

GGP—Gateway-to-Gateway Protocol. GGP is a TCP/IP protocol that transfers routing information between gateways. See also *G.711 A-Law, TCP/IP*.

GHz—Gigahertz. GHz is a unit of wave frequency equal to one thousand million hertz (1,000,000,000 Hz). In some computers, microprocessor clock speed is measured in GHz. Personal computer clock speeds are generally a few tenths of a GHz, but are increasing toward 1 GHz.

glare—A signal that the switch sends when you attempt to place an outgoing call and answer an incoming call simultaneously.

global hunt-group number—A phone number that spans all the T1s of all the MAX units in a stack. The telephone company has set up the global hunt group to distribute incoming calls equally among the MAX units.

global IP address pool—A pool used by several MAX units for dynamically allocating IP addresses to dial-in clients. By default, each MAX handles dynamic IP address allocation from a pool of addresses individually assigned to it. However, you can also set up RADIUS to allocate IP addresses from a global pool that many units share. To do so, you must install RADIPAD, the central manager for global IP address pools on a network. Although multiple hosts can run the RADIUS daemon, only one host on the network should run RADIPAD. See also *dynamic IP*, *IP address pool*, *RADIPAD*, *RADIUS*, *RADIUS daemon*.

Global System for Mobile Communication—See *GSM*.

GloBanD—A European data service consisting of a single circuit whose bandwidth is a multiple of 64 Kbps. This circuit consists of one or more B channels. For example, if a caller requests 512-Kbps service, the line uses eight B channels to supply the requested bandwidth. GloBanD service is available over T1 PRI lines only. It differs from MultiRate in being an overlay network, rather than an integral part of the worldwide switched digital infrastructure. See also *bandwidth*, *B channel*, *MultiRate*, *T1 PRI line.*

GMT—Greenwich Mean Time. This term has been changed to *Coordinated Universal Time (UTC)*. See *UTC*.

GRE—Generic Routing Encapsulation. GRE provides a simple, general-purpose mechanism for encapsulating an arbitrary Network-layer protocol in another arbitrary Network-layer protocol. When a system needs to route data, it first encapsulates the information in a GRE packet. The system then encapsulates the GRE packet in a protocol supported by the network and forwards the packet to its destination.

Greenwich Mean Time—This term has been changed to *Coordinated Universal Time (UTC)*. See *UTC*.

ground start signaling—A signaling method in which the Customer Premises Equipment (CPE) transmits an off-hook condition by creating a zero-voltage condition. Compare with *loop start signaling*, *wink-start signaling*.

GSM—Global System for Mobile Communication. GSM is the most commonly used digital wireless telephone technology. It performs analog-to-digital (A-D) conversion, compressing data and transmitting it on a channel with two other data streams, each in its own time slot. Compare with *CDMA*. See also *wireless technology*.

GSM 1900—Also known as *PCS 1900* or *DCS 1900*, one of the three Personal Communications Services (PCS) technologies in North America. GSM is the only one that provides data services and allows movement between North America and Europe. Omnipoint, Pacific Bell, BellSouth, Sprint Spectrum, Microcell, Western Wireless, Powertel, and Aerial all support GSM 1900.

H

H0 channel—In the Switched-384 data service, a circuit consisting of 6 B channels, or 384 Kbps. See also *B channel, Switched-384.*

H0 data service—See *Switched-384.*

H11 channel—In the Switched-1536 data service, a circuit consisting of 24 B channels, or 1536 Kbps. See also *B channel, Switched-1536.*

H11 data service—See *Switched-1536.*

half duplex—A type of communications configuration in which data can be transmitted in only one direction at a time. Compare with *full duplex*.

handshaking—The process of exchanging signaling information between two communications devices in order to establish the manner and speed of data transmission. You can use either hardware handshaking or software handshaking. See also *hardware handshaking, software handshaking*.

hardware address—An address assigned by the hardware manufacturer and unique to a device.

hardware flow control—A method of flow control that uses separate wires in the modem cable to signal stop and start requests between two directly connected systems. Compare with *software flow control*. See also *flow control*.

hardware interface—A hardware link between two devices. A hardware interface has electrical, physical, and functional specifications that determine how two devices communicate. An electrical specification defines the characteristics of the electrical signals. A physical specification might define the number of pins and wires required, and the order in which the pins and wires are laid out. The functional specification instructs the hardware on how to interpret the electrical signals. Examples of commonly used hardware interfaces are RS-232 and V.24. See also *interface, RS-232, V.24.*

hardware handshaking—A method of synchronizing data transmissions by using the Request To Send (RTS) and Clear To Send (CTS) pins on a wire. Compare with *software handshaking*. See also *CTS, handshaking, RTS*.

HDLC—High-Level Data Link Control. HDLC is a synchronous, bit-oriented Data Link layer protocol for data transmission. Frame Relay is an example of an HDLC-based packet protocol. HDLC offers half- or full-duplex communications over circuit- or packet-switched networks, allows point-to-point and multipoint configurations, and provides transmission over both wires and wireless media. See also *circuit switching, Data Link layer, Frame Relay, full duplex, half duplex, multipoint link, packet switching, point-to-point link, wireless technology*.

Hello interval—For an Open Shortest Path First (OSPF) router, the number of seconds between sending OSPF Hello packets on the interface. See also *Hello packet*, *OSPF*.

Hello packet—A packet that an Open Shortest Path First (OSPF) router uses to recognize when a router is down. After it stops receiving the router's Hello packets, the MAX waits a specified period of time before declaring its neighboring routers down. See also *Hello interval*, *OSPF*.

High-Level Data Link Control—See *HDLC*.

High-Speed Serial Interface—See *HSSI*.

HMP—Host Monitoring Protocol. HMP is a protocol for collecting information from hosts on various networks, including servers, workstations, switches, and gateways. Using HMP, a device can monitor Internet hosts as well as hosts on a private network.

Home Agent—An Ascend unit that represents the terminating part of the Ascend Tunnel Management Protocol (ATMP) tunnel. It must be able to communicate with the Home Network directly, through another router, or across a nailed-up WAN connection. See also *ATMP*, *Home Network*, *nailed-up circuit*, *router*.

Home Network—A private corporate network in an Ascend Tunnel Management Protocol (ATMP) configuration. A private network is one that cannot communicate directly on the Internet, such as an IP network with an unregistered network number. See also *ATMP*, *IP network*, *IPX network*, *IP network number*.

home server proxy—See *SAP home server proxy*.

hop—A single message or packet transmission between host and a router, or between two routers. See also *hop count*, *host*, *router*.

hop count—The number of routers through which a packet passes to get from its source to its destination. See also *hop*, *host*, *router*.

host—A computer on a network, also called a *node* or a *station*.

Host/6 card—A MAX card that supports up to 32 online channels. You can install a maximum of two Host/6 cards in the MAX.

host address—The IP address of a node on a network. See also *IP address*.

Host/BRI module—On the MAX, a module that enables the unit to act as an ISDN switch and provide ISDN BRI connections for applications such as videoconferencing and distance learning. Each Host/BRI module supports up to eight local ISDN BRI lines. The device terminating each line can be a Pipeline 50 (or any other ISDN BRI device) on its own local Ethernet segment, or a desktop video device with its own ISDN BRI line and built-in Terminal Adapter (TA). To the terminating equipment, such as a telephone, a Host/BRI line appears to be an AT&T switch. See also *Host/BRI port*, *ISDN BRI line*.

Host/BRI port—On the MAX, a DCE port that provides a point-to-point ISDN BRI connection with another device. The Host/BRI module has eight Host/BRI ports. From the point of view of the MAX, pins 3 and 6 receive on the Host/BRI interface, while pins 4 and 5 transmit on the Host/BRI interface. See also *Net/BRI port*.

Host/Dual module—On the MAX, a module that manipulates the base system's two AIM ports.

Host Monitoring Protocol—See *HMP*.

host number—The portion of an IP address that denotes an individual node on a network. The class of an IP address determines which portion of the address belongs to the network number and which portion belongs to the host number. See also *IP address*, *IP network number*.

host route—An IP address with a subnet mask of 255.255.255.255, representing a single host rather than a remote router. A host route requires a static IP address. Figure 21 shows a sample connection in which a dial-in host with an ISDN modem card calls into the MAX and requires a static address for a host route.

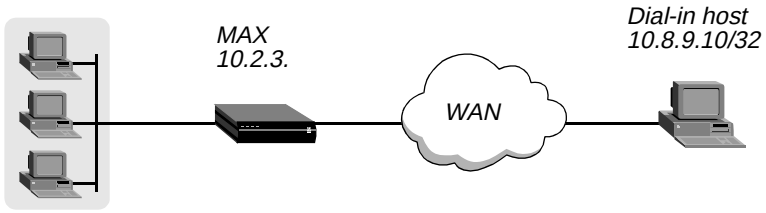

Figure 21. Dial-in host requiring a static IP address (a host route)

See also *host*, *route*, *router*, *subnet mask*.

host-route connection—A connection that enables the dial-in host to keep its own IP address when logging into the MAX IP network. See also *host route*.

host-side port—A port on the device that terminates the WAN circuit for an incoming call to the MAX.

HSSI—High-Speed Serial Interface. HSSI is a serial interface that operates at speeds of up to 52 Mbps, and at distances of up to 50 feet. It is similar to the RS-232 and V.35 serial interfaces, but operates at a higher speed. See also *RS-232*, *V.35*.

HTTP—Hypertext Transfer Protocol. HTTP enables Internet users to request, receive, and provide documents on the World Wide Web.

hub—A device that serves as a termination point for multiple hosts, sending signals onto the proper paths. Typically, a hub contains four to eight connectors. In addition to providing connectors for hosts, many hubs include connectors that you can use to link one hub to another.

hunt group—A group of channels that share the same phone number. When a call comes in using the phone number assigned to the hunt group, the switch hunts for an available channel in the group. See also *channel*, *switch.*

hybrid LAN—A network in which some links are capable of sending and receiving analog signals, while others handle digital signals. See also *analog signal*, *digital signal*.

Hypertext Transfer Protocol—See *HTTP*.

I

IAB—Internet Architecture Board. Part of the Internet Society, the IAB oversees technical innovation on the Internet. The IAB supervises a number of committees, including the Internet Engineering Task Force (IETF), the Internet Research Task Force (IRTF), the Internet Assigned Numbers Authority (IANA), and the Internet Registry. See also *IANA*, *IETF*, *Internet Registry*, *IRTF*.

IAD—Integrated Access Device. An IAD is a device that plugs into a computer by means of an Ethernet connection, and to an analog phone, fax machine, or answering machine by means of an analog port. The IAD supports all the communication devices in a home office using a single ISDN line. The ISDN line replaces multiple analog lines while providing improved speed and throughput. The IAD also supports bridging and dynamic bandwidth management. A Pipeline 85 is an example of an IAD. See also *ISDN line*.

IANA—Internet Assigned Numbers Authority. The IANA is the part of the Internet Society responsible for assigning IP addresses to new Points of Presence (POPs). See also *Internet Society*, *POP*.

ICMP—Internet Control Message Protocol. ICMP is an error-reporting mechanism integral to the TCP/IP protocol suite. Gateways and hosts use ICMP to send reports of datagram problems to the sender. ICMP also includes an echo request/reply function that tests whether a destination is available and responding. See also *G.711 A-Law*, *host*, *TCP/IP*.

ICMP Redirect packet—A packet that instructs the receiver to override a setting in its routing table. A router can use an ICMP Redirect packet to tell a host that it is sending packets to the wrong router, and to inform the host of the correct route. However, a forged ICMP Redirect packet can alter the host's routing table and compromise the security of the network. For this reason, many firewall builders prohibit ICMP traffic from their networks. See also *DoS attack*, *firewall*, *ICMP*.

idle connection—A connection between endpoints in which no data is being transmitted.

idle disconnect—A disconnect that occurs when no data is transmitted on a link for a specified period of time. See also *idle timer*.

idle limit—The value specified for the Ascend Tunnel Management Protocol (ATMP) inactivity timer. See also *inactivity timer*.

idle timer—A timer that measures how long a session can remain idle before the MAX disconnects it. By default, any traffic across an active connection resets the connection's idle timer. When you apply a call filter, its forwarding action determines which packets can initiate a connection or reset a session's timer. When a session's idle timer expires, the MAX terminates the session. The idle timer is set to 120 seconds by default, so if a connection is inactive for two minutes, the MAX terminates the connection. Compare with *inactivity timer*. See also *call filter*.

IDRP—Inter-Domain Routing Protocol. IDRP is an International Standards Organization (ISO) protocol for routing packets between disparate administrative domains. It is based on the Border Gateway Protocol (BGP). See also *BGP*, *ISO*.

IDSL card—ISDN Digital Subscriber Line card. An IDSL card supports eight ports for incoming and outgoing data and voice transmissions at speeds of up to 128 Kbps over 18,000 feet (with single copper-pair wiring) and over longer distances with Digital Loop Carriers (DLCs) and U-loop repeaters. The MAX receives outgoing call requests from attached Terminal Equipment (TE) and routes voice calls to the Public Switched Telephone Network (PSTN) over a T1 or PRI line. For the MAX to support outgoing voice calls, the connected TE must send digits to the MAX using Q.931 en-bloc dialing. The MAX receives incoming voice calls and routes them to a TE connected to an IDSL card by means of Dialed Number Identification Service (DNIS). See also *DLC*, *DNIS*, *IDSL port*, *PSTN*, *T1 line*, *T1 PRI line*, *TE*.

IDSL port—ISDN Digital Subscriber Line port. An IDSL port is a DCE port that provides point-to-point IDSL connections between the MAX and another device. The IDSL card has eight IDSL ports. From the point of view of the MAX, pins 3 and 6 receive on the IDSL port, while pins 4 and 5 transmit on the IDSL port. See also *IDSL card*.

IEC—Interexchange Carrier. An IEC is a type of telephone service that provides long-distance links between local telephone companies. Well-known IECs include AT&T, MCI, and Sprint. Compare with *LEC*.

IEEE—Institute of Electrical and Electronics Engineers. The IEEE is an organization that maintains the standards for 10Base-T and other communications specifications. See also *10Base-T*.

IETF—Internet Engineering Task Force. Responsible for developing the TCP/IP protocol suite, the IETF operates under the auspices of the Internet Society's Internet Architecture Board. See also *Internet Society*.

I-frame—Information frame. An I-frame transports data over an X.25 access link. Compare with *general frame*, *Supervisory frame*. See also *X.25*.

IGMP—Internet Group Management Protocol. IGMP is a protocol implemented by multicast clients and routers. The MAX responds as a client to IGMP packets it receives from a Multicast Backbone (MBONE) router. The packets may use IGMP version-1, IGMP version-2, or IGMP Multicast Trace (MTRACE). MAX clients wanting MBONE service must implement IGMP. See also *MBONE*, *multicast*, *multicast network*, *router*.

IGP—Interior Gateway Protocol. IGP transmits routing information internal to a network. See also *routing*.

immediate mode—A terminal-server access mode in which the terminal server does not display the command-line prompt or a menu of hosts, but immediately directs a dial-in user to a designated host by means of TCP, Rlogin, or Telnet. When you use Telnet to initiate the connection to the host, you can configure the terminal server to pass the call to the host before authentication. In this case, the responsibility for authentication belongs to the Telnet host. See also *Rlogin*, *TCP*, *Telnet*.

immediate modem service—A feature that enables terminal-server users to have direct access to a particular Telnet port on the MAX for modem dialout without using the terminal-server interface. See also *modem dialout*.

inactivity timer—Under the Ascend Tunnel Management Protocol (ATMP), a timer that determines the number of minutes that a Home Agent maintains an idle tunnel before disconnecting it. Under X.25/IP, the number of seconds the MAX enables a connection to remain inactive before dropping the Virtual Circuit (VC). Compare with *idle timer*. See also *VC*.

inband signaling—A type of signaling in which a line uses 8 Kbps of each 64-Kbps channel for WAN synchronization and signaling. The remaining 56 Kbps handle the transmission of user data. When a line is configured for inband signaling, the MAX does not receive bearer-capability information from the carrier. Therefore, it does not know whether a call uses voice service or digital service. For call-routing purposes, the MAX assumes that all calls on an inband-signaling line are digital. Another term for inband signaling is *robbed-bit signaling*. Robbed-bit refers to the 8 Kbps of each channel used for signaling. Compare with *ISDN D-channel signaling*.

incoming call—A call the MAX receives from a remote user or device.

Information frame—See *I-frame*.

input filter—A filter applied to an incoming packet. See also *filter*, *packet filter*.

Input/Output—See *I/O*.

Institute of Electrical and Electronics Engineers—See *IEEE*.

Integrated Access Device—See *IAD*.

Integrated Services Digital Network—See *ISDN*.

Integrated Small Digital Exchange—See *ISDX*.

Inter-Domain Routing Protocol—See *IDRP*.

Interexchange Carrier—See *IEC*.

interface—A connection between two devices, programs, or program elements. See also *hardware interface*.

interface-based routing—An IP-routing method in which each physical or logical interface on the unit has its own IP address. The interface becomes a numbered interface. Reasons for using numbered interfaces include troubleshooting nailed-up point-to-point connections and forcing routing decisions between two links going to the same final destination. More generally, interface-based routing allows the MAX to operate more as a multihomed host behaves.

You can configure each link as numbered (interface-based) or unnumbered (system-based). If no interfaces are numbered, the MAX operates as a purely system-based router. Compare with *system-based routing, unnumbered interface*. See also *IP routing, multihomed host, numbered interface, point-to-point link.*

interface cost—See *link-state metric*.

Interior Gateway Protocol—See *IGP*.

internal Link-State Advertisement—See *internal LSA*.

internal LSA—Internal Link-State Advertisement. An internal LSA is exchanged by Open Shortest Path First (OSPF) routers within a single Autonomous System (AS). Compare with *external LSA*. See also *LSA*.

International Standards Organization—See *ISO*.

International Telecommunication Union–Telecommunication Standardization Sector—See *ITU-T*.

internet—A series of networks connected by bridges, gateways, or routers. An internet is also called an *internetwork*. See also *G.711 A-Law*, *router*.

Internet—The complex of WANs joining government, university, corporate and private computers in a vast web of network interconnection.

Internet Architecture Board—See *IAB*.

Internet Assigned Numbers Authority—See *IANA*.

Internet Control Message Protocol—See *ICMP*.

Internet Engineering Task Force—See *IETF*.

Internet Group Management Protocol—See *IGMP*.

Internet Network Information Center—See *InterNIC*.

Internet Protocol—See *IP*.

Internet Protocol Control Protocol—See *IPCP*.

Internet Registry—A branch of the Internet Society's Internet Architecture Board (IAB). The Internet Registry supervises the Domain Name System (DNS). See also *DNS*, *IAB*, *Internet Society*.

Internet Reliable Transaction Protocol—See *IRTP*.

Internet Service Provider—See *ISP*.

Internet Society—An international nonprofit organization that focuses on Internet standards, education, and policy issues. Founded in 1992 and located in Reston, Virginia, the Internet Society oversees the Internet Architecture Board (IAB), which in turn oversees a number of committees, including the Internet Engineering Task Force (IETF), the Internet Research Task Force (IRTF), the Internet Assigned Numbers Authority (IANA), and the Internet Registry. See also *IAB*, *IANA*, *IETF*, *Internet Registry*, *IRTF*.

internetwork—See *internet*.

Internetwork Packet Exchange—See *IPX*.

InterNIC—Internet Network Information Center. InterNIC is an organization that provides Internet information services, oversees the registration of Internet addresses and Domain Name System (DNS) names, assigns RFC numbers, and assists users in gaining access to the Internet. See also *DNS*, *RFC*.

interoperability—Compatibility with the devices and services of multiple vendors. Interoperable devices can be integrated into a network containing a wide range of vendor products. Interoperability is a significant factor among expansion considerations, because any device must have the versatility to function in an expanding network structure. The technical elements of interoperability may include a bundle of protocols and a flexible architecture to accommodate upgrades. A remote-access server should include capabilities such as translation, encapsulation, and filtering.

intra-area route—A route imported into the Open Shortest Path First (OSPF) database from within the router's area. Compare with *external route*. See also *area*, *OSPF*, *route*, *router*.

Intragy—An Ascend enterprise access solution that provides corporations with the multiprotocol support needed for flexible access to a corporate network. Intragy combines IntragyCentral (a network-access solution) with IntragyAccess (cross-platform desktop client software). Together, these components provide complete network access for all users of the corporate network—LAN users, telecommuters, remote users, and mobile workers. Multiprotocol routing, Ethernet bridging, desktop client dialout, and desktop client software combine to deliver open access to centralized corporate resources.

inverse multiplexer—Equipment that performs inverse multiplexing at each end of a connection. An inverse multiplexer is also known as an *inverse mux*. See also *inverse multiplexing*.

inverse multiplexing—A method of combining individually dialed channels into a single, higher-speed data stream. Each end of the connection uses an inverse multiplexer (also called an *inverse mux*).

For example, suppose one site has three ISDN BRI lines connected to an inverse mux and another site has a T1 line connected to an inverse mux. The user at the first site can place a 336-Kbps call to the second site using inverse multiplexing. Because each BRI line has two 64-Kbps channels (with 56 Kbps reserved for data on each channel), the inverse mux places six individual calls over Switched-56 services to the answering T1-based inverse mux. The two inverse muxes combine the six calls into a single data stream at 336 Kbps (6X56 Kbps). See also *circuit-level inverse multiplexing*, *inverse multiplexer*, *packet-level inverse multiplexing*, *T1 line*.

inverse mux—An inverse multiplexer.

I/O—Input/Output. A manner of describing a process, program, or device that transfers information to or from a computer. Common I/O devices are hard disks, printers, keyboards, diskettes, and CD-ROMs.

IP—Internet Protocol. IP provides connectionless, non-guaranteed transmission of Transport-layer data packets. IP fragments packets, allowing them to take different paths across the WAN, and then reassembles them into the proper order at their destination. See also *Transport layer*.

IP address—An address that uniquely identifies each host on a network or internet. An IP address has a length of 32 bits, and is divided into four 8-bit parts, each separated by a period, as in 149.122.3.30. This kind of notation is called *dotted decimal notation*. Each part can consist of a number between 1 and 255.

An IP address consists of a network number and a host number. IP addresses come in three types: Class A, Class B, and Class C. The class of an IP address determines which portion of the address belongs to the network number and which portion belongs to the host number. The first bits of the IP address identify the class. The Internet Network Information Center (InterNIC) determines the type of class assigned a network.

A Class A address starts with 0 as the class identifier, followed by 7 bits for the network number and 24 bits for the host number.Therefore, the first number in dotted decimal form is the network number. The next three numbers make up the host number. For example, in the IP address 127.120.3.8, the network number is 127 and the host number is 120.3.8. This type of address is used by the largest organizations, because this scheme allows for over 16 million different host numbers. However, it also limits network numbers to a total of 128.

A Class B address starts with binary 10 as the class identifier, followed by 14 bits for the network number and 16 bits for the host number. Therefore, the first two dotted decimal numbers comprise the network number, and the second two dotted decimal numbers comprise the host number. For example, in the IP address 147.14.86.24, the network number is 147.14 and the host number is 86.24. More network numbers are available than in a Class A address, but fewer hosts (approximately 65,000).

A Class C address starts with binary 110 as the class identifier, followed by 21 bits for the network number and 9 bits for the host number. Therefore, the first three dotted decimal numbers comprise the network number, and the last dotted decimal number comprises the host number. For example, in the IP address 225.135.38.42, the network number is 225.135.38 and the host number is 42. Many network numbers are available, but only 254 hosts per network number. The numbers 0 and 255 are reserved.

You can tell the type of class an IP address falls into by looking at the first 8-bit portion of the dotted decimal form of the address. Class A addresses begin with a number between 0 and 127. Class B addresses begin with a number between 128 and 223. Class C addresses begin with a number between 192 and 233. In addition to an IP address, you can use a symbolic name provided by Domain Name System (DNS) to designate an Internet address. See also *DNS*, *dotted decimal notation*, *host number*, *internet*, *InterNIC*, *IP*, *network*, *IP network number*.

IP address pool—A pool from which the MAX dynamically allocates an IP address to a calling unit. You can configure up to 10 address pools on the MAX, and up to 50 in RADIUS. You can configure address pools for use only by a specific MAX. See also *dynamic IP*, *IP address*.

IP address spoofing—A way for a remote device to illegally acquire a local address in order to break through a firewall or data filter. Compare with *DHCP spoofing*, *IPX spoofing*, *SPX spoofing*, *watchdog spoofing*.

IPCP—Internet Protocol Control Protocol. IPCP is a protocol for configuring, enabling, and disabling the IP protocol modules on both ends of a point-to-point link. IPCP is tied to PPP, and is activated only when PPP reaches the Network-layer protocol phase. IPCP packets received prior to this phase are discarded. Elements of IPCP include packet encapsulation, code fields, and timeouts. See also *IP*, *Network layer*, *point-to-point link*.

IP Direct—A configuration in which the MAX automatically redirects incoming IP packets to a host you specify on the local IP network. When you specify IP Direct, the MAX bypasses all internal routing and bridging tables, and sends all packets it receives on a connection's WAN interface to the specified IP address. Figure 22 shows an example of the traffic flow for IP Direct.

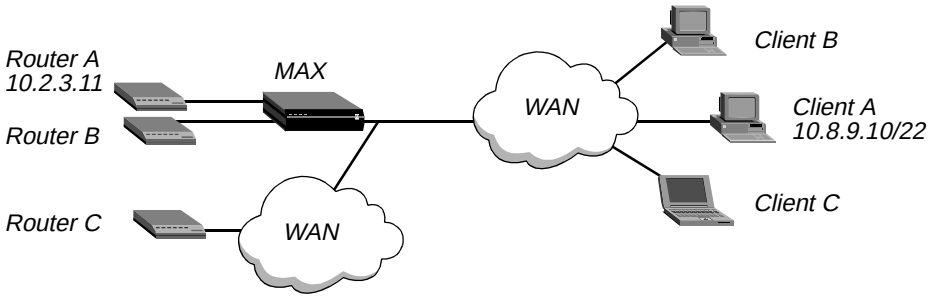

Figure 22. IP Direct connections

The MAX redirects inbound packets from client A to router A, and from client B to router B, on the LAN side of the MAX. From client C, the MAX redirects inbound packets to router C by means of a switched connection.

Packets destined for the clients A, B, or C are routed normally by the MAX, which means that these client connections can *receive* packets from any source, not just from the IP address to which their packets are sent.

See also *IP address*, *IP routing table*.

IP filter—A packet filter that examines fields specific to IP packets. An IP filter focuses on known fields, such as source or destination address and protocol number. It operates on logical information that is relatively easy to obtain. In an IP filter, a number of distinct comparisons occur in a defined order. When a comparison fails, the packet goes on to the next comparison. When a comparison succeeds, the filtering process stops and the MAX applies the forward action in the filter to the packet. Compare with *generic filter*, *IPX filter*. See also *call filter*, *data filter*, *packet filter*.

IP multicast forwarding—See *multicast forwarding.*

IP network—A network that uses the Internet Protocol (IP) to transmit packets at the Transport layer. See also *IP*.

IP network number—The portion of an IP address that denotes the IP network on which a host resides. The class of an IP address determines which portion of the address belongs to the network number and which portion belongs to the host number. See also *host number, IP address.*

IP route—A path from one IP network to another. See also *dynamic route, IP network, multipath route, static route*.

IP router—A device that sends IP packets from a source to a destination by multiple paths. As an IP router, the MAX routes IP packets between its Ethernet interfaces and across any WAN interface configured for IP routing. See also *IP route*, *IP routing.*

IP routing—A method of determining how to forward an IP packet to the proper destination. When acting as an IP router, the MAX routes IP packets between its Ethernet interfaces and across any WAN interface configured for IP routing. Figure 23 shows a MAX that routes IP packets between WAN interfaces and a LAN interface.

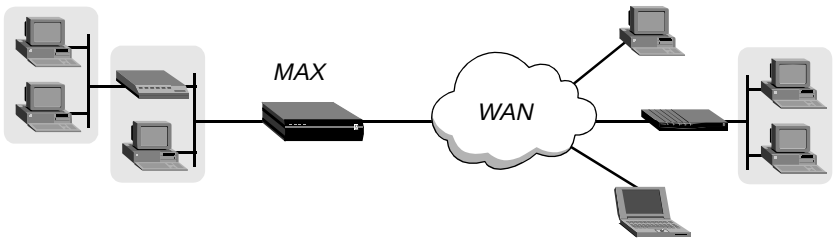

Figure 23. IP-routing configuration

The MAX consults its internal routing table to determine where to forward each IP packet it processes. First, the MAX tries to find a match between the packet's destination address and a Destination field in its routing table. If it finds a match, it brings up the required connection (if necessary) to reach the next-hop router specified for that route, and forwards the packet.

If it does not find a match for the packet's destination address, it looks for a default route (destination address 0.0.0.0). If it finds a default route, it brings up the required connection (if necessary) and forwards the packet. If the routing table has no default route, and no route that matches a packet's destination address, the MAX drops the packet. See also *default route*, *hop*, *IP route*, *IP router*, *IP routing table*.

IP routing table—A table that contains information about how to forward IP packets. On the MAX, the IP routing table looks like the following:

```
Destination        Gateway       IF       Flg    Pref  Metric Use
Age
0.0.0.0/0          206.65.212.1  ie0      SG      100   1     4891
48630
10.0.0.0/24        11.168.6.249  ie1-12-1 RGT    100   3        0
9236
10.0.100.0/24      11.168.6.86   ie1-12-1 RGT    100   2        0
48601
10.0.200.0/24      11.168.6.86   ie1-12-1 RGT    100   2        0
48601
10.122.72.0/24     -             ie1-12-2 C        0   0     3141
48630
10.122.72.1/32     -             lo0      CP       0   0        0
48630
10.122.73.0/24     -             ie1-12-3 C        0   0     3140
48630
10.122.73.1/32     -             lo0      CP       0   0        0
48630
10.122.74.1/32     -             lo0      CP       0   0        0
48630
10.122.99.0/24     10.122.99.1   wan4     SG      100   7        0
48630
10.122.99.1/32     10.122.99.1   wan4     S       100   7        1
48630
```

The fields in the routing table contain the following information:

Field	Description
`Destination`	The route's target address. To send a packet to this address, the MAX uses this route. If the target address appears more than once in the routing table, the MAX uses the most specific route (having the largest subnet mask) that matches that address.
`Gateway`	The next hop router that can forward packets to the given destination. Direct routes (without a gateway) show a hyphen in this column.
`IF`	The name of the interface through which to send packets over this route: • `ie0` or `ie[`*`shelf`*`]-[`*`slot`*`]-[`*`item`*`]` is an Ethernet interface. • `lo0` is the loopback interface. • `rj0` is the reject interface, used in network summarization. • `bh0` is the blackhole interface, used in network summarization. • `wan`*`N`* is a WAN connection, entered as it becomes active. • `wanabe` indicates an inactive RADIUS dialout profile. • `local` indicates a single route targeted at the local machine. • `mcast` indicates a route to a virtual device. The route encapsulates the multicast forwarder for the entire class D address space.
`Flg`	One or more of the following flags: • `C`—a directly connected route, such as Ethernet • `I`—an ICMP redirect dynamic route • `N`—placed in the table via SNMP MIB II • `O`—A route learned from OSPF • `R`—a route learned from RIP • `r`—a transient RADIUS-like route • `S` —a static route • `?`—a route of unknown origin, which indicates an error • `G`—an indirect route via a gateway • `P`—a private route • `T`—a temporary route • `M`—a multipath route • `*`—a backup static route for a transient RADIUS-like route
`Pref`	The preference value. See the description of the Preference parameter for information about defaults for route preferences.
`Metric`	A RIP-style metric for the route, with a range of 0-16. Routes learned from OSPF show a RIP metric of 10. OSPF cost-infinity routes show a RIP metric of 16.
`Use`	A count of the number of times the route was referenced since it was created. (Many of these references are internal, so this is not a count of the number of packets sent over this route.)
`Age`	The age of the route in seconds. RIP and ICMP entries are aged once every 10 seconds.

See also *direct route, dynamic route, G.711 A-Law, hop, IP route, IP router, metric, multipath route, OSPF, preference, RIP, static route*.

IP subnet—A portion of an IP network. IP subnetting is a way to subdivide a network into smaller networks, resulting in a greater number of hosts on a network associated a single IP network number. An IP address that uses a subnet has three elements: network, subnet, and host. You identify a subnet by combining an address with a subnet mask. For example, in the address 195.112.56.75/14, /14 is the subnet mask. See also *host number, IP address, IP network number, subnet mask*.

IP switch—A device that can determine the destination of large volumes of incoming IP packets and send them to the appropriate outgoing ports at high speeds. An IP switch is a high-performance device designed for high-volume, large-scale public and private backbone applications. See also *switch*.

IPX—Internetwork Packet Exchange. IPX is Novell's connectionless Network-layer protocol. Derived from XNS' Internetwork Datagram Protocol (IDP), IPX performs addressing and routing functions. At the server, IPX passes outgoing datagrams to the network interface software. At the packet's destination, IPX passes the data to upper-layer processes. Along an IPX route, intermediate devices use IPX to route packets to their destinations. When routing, IPX relies on information supplied by the Routing Information Protocol (RIP). See also *IPX network, IPX route, IPX routing, IPX server, RIP*.

IPX bridging—At the Data Link layer, a way of passing IPX packets between networks. See also *Data Link layer, IPX network*.

IPX client—A user or device that gains access to the services of an IPX server. See also *IPX server*.

IPXCP—Internet Packet Exchange Control Protocol. IPXCP is a protocol for configuring, enabling, and disabling the IPX protocol modules on both ends of a point-to-point link. IPXCP is tied to PPP, and is activated only when PPP reaches the Network-layer protocol phase. IPXCP packets received prior to this phase are discarded. Elements of IPXCP include packet encapsulation, code fields, and timeouts. See also *IPX, point-to-point link*.

IPX filter—A packet filter that examines fields specific to IPX packets. An IPX filter focuses on known fields, such as source or destination address. It operates on logical information that is relatively easy to obtain. In an IPX filter, a number of distinct comparisons occur in a defined order. When a comparison fails, the packet goes on to the next comparison. When a comparison succeeds, the filtering process stops and the MAX applies the forward action in the filter to the packet. You can configure an IPX filter as either a call filter or a data filter. Compare with *generic filter, IP filter*. See also *call filter, data filter, packet filter*.

IPX frame—The type of packet frame used by an IPX server. An IPX frame can follow the IEEE 802.2, IEEE 802.3, SubNetwork Access Protocol (SNAP), or Ethernet II protocol specification for the Media Access Control (MAC) header. See also *802.2, 802.3, Ethernet II, IPX server, MAC, SNAP*.

IPX Nearest Server Query—In an Ascend Tunnel Management Protocol (ATMP) configuration, a message sent to a Home Agent from a Mobile Client, asking whether the Home Agent knows about a server on the Home Network. See also *ATMP, Home Agent, Home Network, Mobile Client*.

IPX network—A network consisting of one or more IPX servers and IPX clients. See also *IPX client, IPX server, virtual IPX network.*

IPX network number—The portion of an IPX address that denotes the IPX network on which a node resides. If the MAX is routing IPX and there are other IPX servers on the LAN interface, the IPX network number assigned to the MAX for that interface must be consistent with the number in use by the other routers.

If you do not specify an IPX network number in the MAX configuration, the MAX learns its network number from another router on the interface or from the Routing Information Protocol (RIP) packets received from the IPX router. If you specify an IPX network number in the MAX configuration, the MAX becomes a seed router, and other routers can learn their network number from the MAX. See also *IPX router*, *seed router*.

IPX route—A path from one IPX network to another. See also *IPX network, IPX router.*

IPX router—A device that sends IPX packets from a source to a destination by various paths. See also *IPX route.*

IPX routing—A method of sending IPX packets from a source to a destination at the Network layer. A MAX configured for IPX routing enables NetWare clients and distributed Novell networks to use NetWare services across the WAN. The NetWare version must be 3.11 or later. Figure 24 shows a MAX that routes IPX between WAN interfaces and a local Novell network.

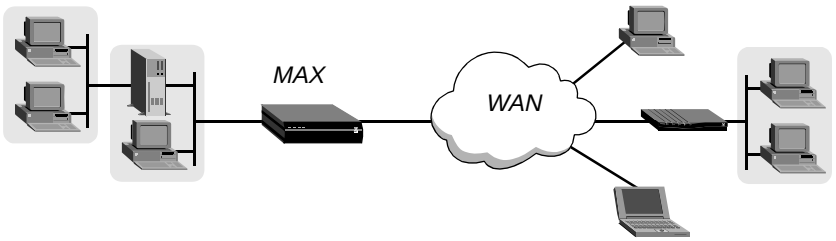

Figure 24. IPX-routing configuration

Ascend has optimized IPX routing for the WAN. The MAX incorporates changes to standard IPX behavior with regard to IPX Service Advertising Protocol (SAP), IPX RIP, and IPXWAN negotiation. In addition, Ascend has added IPX extensions that enable the MAX to operate as clients expect for NetWare LANs. These extensions include creating a virtual network for dial-in IPX clients, accepting or rejecting RIP and SAP updates, bringing up connections in response to a SAP query, creating static IPX routes, and defining SAP filters.

See also *dial query*, *IPX*, *IPX network*, *IPX route*, *IPX server*, *IPXWAN*, *RIP*, *SAP*, *SAP filter*, *static IPX route*, *virtual IPX network.*

IPX SAP—See *SAP*.

IPX SAP filter—See *SAP filter*.

IPX server—A server that runs the NetWare operating system, manages network resources, and communicates with IPX clients. See also *IPX, IPX client.*

IPX spoofing—A procedure that enables a device to mimic a legitimate network host and gain access to data within a private IPX network. Spoofing can lead to severe security breaches and damage the integrity of a company's operations. Compare with *DHCP spoofing*, *IP address spoofing*, *SPX spoofing*, *watchdog spoofing*. See also *IPX network.*

IPX Type 20—A type of packet that applications such as NetBIOS over IPX use to broadcast names over a network. By default, these broadcasts are not propagated over routed links (as Novell recommends) and are not forwarded over links that have less than 1 Mbps throughput. However, if you are using an application such as NetBIOS over IPX, which requires these packets in order to operate, you can enable the router to propagate IPX Type 20 packets over a LAN interface.

IPXWAN—The WAN version of NetWare's IPX protocol. The MAX supports the IPXWAN protocol, which is essential for communicating with Novell software that supports dial-in connections, and with the Multi-Protocol Router. For full specifications of the IPXWAN protocol, see RFC 1634 and *NetWare Link Services Protocol Specification—IPX WAN Version 2*.

When an IPX connection is brought up between two Ascend units, all options are negotiated during the IPXCP phase. IPXWAN negotiation never takes place between two Ascend units, because neither unit initiates the negotiation process by sending out an IPXWAN Timer_Request packet.

Connections with non-Ascend devices that use Novell software operating over PPP do not negotiate options during the IPXCP phase, so all options are negotiated during the IPXWAN phase of link establishment. The remote device sends an IPXWAN Timer_Request packet, which triggers IPXWAN negotiation in the MAX. The devices compare internal network numbers and assign the slave role to the unit with the lower number. The other unit becomes the master of this link for the duration of the IPXWAN negotiation. The slave unit returns an IPXWAN Timer_Response packet, and the master unit initiates an exchange of information about the final router configuration. The MAX supports the following routing options:

- Ascend Routing (Unnumbered RIP/SAP without aging).
- Novell Routing (Unnumbered RIP/SAP with aging).
- None (The peer is a dial-in client. No RIP/SAP except on request. Network and node numbers may be assigned.)

Header compression is rejected as a routing option. After IPXWAN negotiation is complete, transmissions of IPX packets use the negotiated routing option.

See also *IPX*, *IPXCP*, *RIP*, *SAP*.

IRTF—Internet Research Task Force. The IRTF is overseen by the Internet Society's Internet Architecture Board (IAB) and is composed of a number of focused, long-term research groups. These groups are concerned with issues related to Internet protocols, architecture, and technology. See also *IAB*, *Internet Society.*

IRTP—Internet Reliable Transaction Protocol. IRTP is a full-duplex, transaction-oriented, host-to-host protocol providing reliable, sequenced delivery of data packets and a constant connection between two hosts. Unlike the User Datagram Protocol IRTP provides reliable, sequenced delivery of packets; unlike the Transaction Control Protocol (TCP), IRTP sequencing takes place on a packet-by-packet basis, and only one connection is defined between Internet addresses. See also *TCP*, *UART*.

ISDN—Integrated Services Digital Network. ISDN is a telecommunications architecture capable of sending voice, data, and video in digital form on a digital line. It can support bandwidths of up to 2 Mbps, and uses a single digital line for telephone, fax, computer, and video communications. ISDN supports circuit-switched and Frame Relay connections. See also *circuit switching*, *digital data*, *E1 PRI line*, *Frame Relay*, *ISDN BRI line*, *T1 PRI line*.

ISDN Basic Rate Interface line—See *ISDN BRI line*.

ISDN BRI line—ISDN Basic Rate Interface line. An ISDN line uses two B channels for user data, and one 16-Kbps D channel for ISDN D-channel signaling. Both B channels can be switched, both channels can be nailed up, or one channel can be switched and the other nailed up. An ISDN BRI line can connect to standard voice service, the Switched-56 data service, or the Switched-64 data service. Compare with *E1 PRI line*, *T1 PRI line*. See also *B channel*, *D channel*, *ISDN D-channel signaling*, *nailed-up channel*, *Switched-56*, *Switched-64*, *switched channel*.

ISDN BRI network-interface card—A MAX card that contains eight ISDN BRI ports. You can install a maximum of four ISDN BRI network-interface cards in the MAX.

ISDN BRI terminal-interface card—A MAX card that contains eight ISDN BRI ports. You can install a maximum of four ISDN BRI terminal-interface cards in the MAX.

ISDN Call Setup message—A message that an ISDN device sends to the network when you make a call. The Call Setup message requests a voice or data bearer service. The network passes along the information in the Call Setup message until it reaches the party you are calling, and then sends a Call Setup message to the ISDN device receiving the call. The called device routes the call to a telephone or to a computer, depending on whether the Call Setup message specifies voice or data bearer service. If the message specifies a data call, the device signals the computer. If the message specifies a voice call, it rings the telephone. See also *bearer service*, *ISDN*.

ISDN D-channel signaling—A type of signaling in which a D channel handles WAN synchronization and signaling, and the B channels carry the user data. Another term for ISDN D-channel signaling is *out-of-band signaling*. T1 PRI, E1 PRI, and ISDN BRI lines use ISDN D-channel signaling. See also *B channel*, *D channel*, *E1 PRI line*, *ISDN BRI line*, *T1 PRI line*.

ISDN Digital Subscriber Line card—See *IDSL card*.

ISDN Digital Subscriber Line port—See *IDSL port*.

ISDN Disconnect message—A message that initiates the release of an ISDN connection. See also *ISDN*.

ISDN line—A line that uses ISDN D-channel signaling. E1 PRI, ISDN BRI, and T1 PRI are all examples of ISDN lines. See also *E1 PRI line*, *ISDN BRI line*, *ISDN D-channel signaling*, *T1 PRI line*.

ISDN modem—See *V.120 TA*.

ISDN PRI line—See *T1 PRI line*.

ISDN Primary Rate Interface line—See *T1 PRI line*.

ISDN Q.931 Layer 3 SETUP_ACK timer—See *T302 timer*.

ISDN subaddress—See *subaddress*.

ISDX—Integrated Small Digital Exchange. ISDX is a telephone switch manufactured by GEC Plessey Telecom (GPT).

ISLX—A DPNSS switch type. See also *DPNSS*.

island—A group of networks on the Multicast Backbone (MBONE). The islands are connected by tunnels and support IP. See also *MBONE*.

ISO—International Standards Organization. The ISO is an organization devoted to the definition of standards for national and international data communications. The U.S. representative to the ISO is the American National Standards Institute (ANSI). Companies whose products are ISO certified reflect a high quality of consistency and quality.

ISO 9001—The current set of International Standards Organization (ISO) standards. See also *ISO*.

ISO Transport Protocol—A protocol that enables peers to exchange information in discrete units called Transport Protocol Data Units (TPDUs). See also *ISO*.

ISP—Internet Service Provider. An ISP is a company that provides access to the Internet. By establishing Points of Presence (POPs) containing remote-access servers and a suite of user software packages, the ISP acts as a commercial on-ramp to the Internet. Providers typically charge a monthly fee, and supply technical support and advice to customers.

ITU—International Telecommunication Union. Headquartered in Geneva, Switzerland, the ITU is an international organization that enables governments and the private sector to coordinate global telecommunication networks and services.

ITU Annex A—See *Frame Relay Annex A*.

ITU-T—International Telecommunication Union–Telecommunication Standardization Sector. The ITU-T is the committee that replaced the Consultative Committee for International Telegraphy and Telephony (CCITT) on March 1, 1993. The ITU-T is responsible for a wide array of telecommunications and networking standards.

J

Java—An object-oriented programming language developed by Sun Microsystems, Inc. You can use Java to create applets for distribution on the World Wide Web. Java programs run inside a Java-enabled Web browser or inside a Java Virtual Machine (JVM).

Java-based configurator—A utility that guides you through the configuration and management of an Ascend unit by means of a graphical user interface (GUI).

Java Virtual Machine—See *JVM*.

JEDEC—Joint Electronic Device Engineering Council. JEDEC is an organization that creates and supervises industry standards for electronic devices. See also *JEDEC file*.

JEDEC file—A text file containing information for configuring a device. See also *JEDEC*.

Joint Electronic Device Engineering Council—See *JEDEC*.

JVM—Java Virtual Machine. A JVM is an abstract computer that runs compiled Java code. The JVM is *virtual* because it is software that runs on top of a hardware platform and an operating system. All Java programs are compiled for a JVM. See also *Java*.

K

K56flex—A 56-Kbps modem specification developed by Rockwell and Lucent for calls between a digital modem and an analog modem. K56flex allows 56-Kbps data transfers on the downstream portion of a call, and 33.6-Kbps data transfers on the upstream portion. The Series56 Digital Modem module includes the K56flex technology. See also *digital modem*, *modem*, *Series56 Digital Modem module*.

L

L2F—Layer Two Forwarding. L2F is a protocol that permits a system to tunnel the Data Link layer (HDLC or SLIP) frames associated with higher-level protocols (such as PPP). See also *HDLC*, *PPP*, *SLIP*.

L2TP—Layer-2 Tunneling Protocol. L2TP enables you to dial into a local ISP and connect to a private corporate network across the Internet. You dial into a local MAX, configured as an L2TP Access Concentrator (LAC), and establish a PPP connection. Attributes in your RADIUS user profile specify that the MAX, acting as a LAC, establish an L2TP tunnel. The LAC contacts the L2TP Network Server (LNS), which is connected to the private network. The LAC and the LNS establish an L2TP tunnel (by means of UDP), and any traffic your client sends is tunneled to the private network. Once the MAX units establish the tunnel, the client connection has a PPP connection with the LNS, and appears to be directly connected to the private network. See also *LAC*, *LNS*.

L2TP Access Concentrator—See *LAC*.

L2TP Network Server—See *LNS*.

LAC—L2TP Access Concentrator. A LAC performs the following functions:

- Establishes PPP connections with dial-in clients.
- Sends requests to L2TP Network Server (LNS) units requesting creation of tunnels.
- Encapsulates and forwards all traffic from clients to the LNS via the tunnel.
- De-encapsulates traffic received from an established tunnel, and forwards it to the client.
- Sends tunnel-disconnect requests to LNS units when clients disconnect.

See also *L2TP*, *LNS*.

LAN—Local Area Network. A LAN is a network in which two or more computers, located within a limited distance of one another, are connected in order to share files and resources. A PC-based LAN consists of a dedicated server running a network operating system and attached to several workstations. A host-based LAN consists of one or more hosts and terminals. Examples of LAN architectures are Ethernet, ARCnet, Fiber Distributed Data Interface (FDDI), and Token Ring. See also *ARCnet*, *FDDI*, *Token Ring*.

LAN adapter—See *NIC*.

landline telephone communication—A communications method in which a signal is carried over a copper local loop. Compare with *cellular communication.*

LAN-to-LAN modem access—A configuration in which two remote-access devices route or bridge traffic between LANs by means of a dial-up modem connection. See also *dial-up line*.

LAN UTP port—Local Area Network Unshielded Twisted Pair port. A LAN UTP port connects the MAX to a UTP LAN.

LAN/WAN connectivity—The ability to link Local Area Networks (LANs) and Wide Area Networks (WANs). A wide range of tools, from translation protocols to communications features to support services, make a remote-access device like the MAX an effective link between LANs and WANs.

LAP—Link Access Procedure. LAP is a protocol containing a subset of High-Level Data Link Protocol (HDLC) features. In order to maintain compatibility with HDLC, LAP was changed to create LAPB. See also *LAPB*.

LAPB—Link Access Procedure, Balanced. LAPB is a protocol for B channels that use packet-switching mode. See also *B channel, packet switching*.

LAPB T1 timer—On an X.25 link, a timer that specifies the maximum amount of time in seconds the transmitter should wait for an acknowledgment before initiating a recovery procedure. On a transmission line between a user and the network, a particular frame or acknowledgment may be incorrectly transmitted or simply discarded. To keep the transmitter from waiting indefinitely for an acknowledgment, you can specify the maximum amount of time the transmitter should wait. When you choose a value, you must take into account any frame transmission and processing delays you may encounter. In most cases, you should use the default value suggested by the network. See also *LAPB, X.25*.

LAPD—Link Access Procedure, D channel. LAPD is a protocol for the D channel. It provides the mechanism for combining multiple channels into a single logical link, and for monitoring and controlling the flow of data over the B channels. See also *B channel, D channel*.

LAPF—Link Access Procedure, Frame. LAPF is a protocol for Frame-mode bearer services. See also *bearer service*.

LAPM—Link Access Procedure, Modem. LAPM is an error-detection protocol for correcting data communication errors occurring on the link between two modems.

LAT—Local Area Transport. Developed by Digital in 1981, the LAT protocol enables you to connect a number of asynchronous devices, such as terminals, printers, modems, and hosts, on an Ethernet network.

latency—For a communications channel, the amount of time before the channel is available for a transmission; for data transmissions, the amount of time it takes for a packet to reach its destination. The following elements contribute to latency:

- The type of physical media in use.
- Physical interference from noise or other signals.
- Required setup and teardown times.
- Signal interfaces. Ethernet consumes a minimum of 0.3 milliseconds (ms). A 28.8 modem takes about 300 times longer.
- Bottlenecks, such as the 50 ms it takes to move data through a serial port.
- Data conversion, such as the conversion from digital to analog data required by a modem.
- Compression.

Once latency is present, it cannot be optimized. You must remove the cause. To maximize throughput, use the highest bandwidth available. All services go as fast as the medium allows. For example, if the medium is copper, the speed of the electrical signal through the copper does not vary with the type of line in use. A T1 line is considered faster than a single analog line only because its bandwidth is greater.

Layer-2 Tunneling Protocol—See *L2TP*.

Layer Two Forwarding—See *L2F*.

LCN—Logical Channel Number. On an X.25 link, an LCN is a unique number assigned to each Virtual Circuit (VC). See also *VC*, *X.25*.

LCP—Link Control Protocol. LCP sets up, manages, and tears down a connection between two Point-to-Point Protocol (PPP) endpoints. See also *PPP*.

learning bridge—See *transparent bridge*.

leased address—In a Network Address Translation (NAT) for LAN configuration, an IP address offered by a Dynamic Host Configuration Protocol (DHCP) server for a limited duration. See also *DHCP*, *DHCP server*, *IP address*, *NAT for LAN*.

leased circuit—See *nailed-up circuit*.

leased line—See *nailed-up line*.

lease time—As defined by the Dynamic Host Configuration Protocol (DHCP) in a Network Address Translation (NAT) for LAN configuration, the time in which a host is assigned an IP address. If the host renews the address before its lease period expires, the DHCP service reassigns the same address. Plug and Play addresses always expire in 60 seconds. See also *DHCP*, *DHCP server*, *IP address, leased address*, *NAT for LAN*.

LEC—Local Exchange Carrier. An LEC is a local telephone company. See also *IEC*.

Level 2 Window Size—On an X.25 connection, the maximum number of sequentially numbered frames that a given Data Terminal Equipment (DTE)/Data Circuit-Terminating Equipment (DCE) link may have unacknowledged at any given time. See also *DCE*, *DTE*, *X.25*.

line—A physical interface to the WAN. A line consists of one or more channels, each of which can transmit data. See also *channel*.

line buildout—See *buildout*.

Line Quality Monitoring—See *LQM*.

line-side connection—A link that extends from the telephone company's Central Office (CO) to the customer. Line-side connections can be high- or low-bandwidth, digital or analog. Compare with *trunk-side connection*. See also *analog line*, *digital line*.

Line Status windows—On the MAX, windows showing the status of the lines in slot 1. By default, the line status is shown in the top two status windows:

```
|--------------------| |--------------------|
|10-100 1234567890   | |10-200 1234567890   |
| L1/LA nnnnnnnnnn   | | L2/RA ..........   |
|   12345678901234   | |   12345678901234   |
|   nnnnnnnnnnnnnn   | |   ..............   |
|--------------------| l--------------------|
```

Each window displays four lines:

- The first line shows the menu number and column numbers for channels 1–10.
- The second line identifies the line (L1 or L2) and shows a a two-character link-status indicator and one-character channel-status indicator. For example:
  - LA indicates *Link Active* (the line is physically connected).
  - n means the channel is nailed.
  - * indicates a current connection.
  - - means the channels is idle but in service.
  - s means the channel is an active D channel (ISDN only).
- The third line has column headers for channels 11–24.
- The fourth line shows a one-character channel-status indicator for channels 11–24.

See also *status window*.

Line Termination Mode—See *LT mode*.

Link Access Procedure—See *LAP*.

Link Access Procedure, Balanced—See *LAPB*.

Link Access Procedure, D Channel—See *LAPD*.

Link Access Procedure, Frame—See *LAPF*.

Link Access Procedure, Modem—See *LAPM*.

link compression—A process that removes waste and redundancy from the data on a connection, enabling faster throughput. For the MAX to use link compression, both sides must be configured to use the same compression method. You can use Stac compression (an Ascend-modified version of draft 0 of the CCP protocol), Stac-9 compression (the method specified by draft 9 of the Stac LZS compression protocol), or Microsoft Stac compression (the method implemented by Windows 95). See also *CCP*, *slot compression*, *VJ compression*.

Link Control Protocol—See *LCP*.

Link Management Interface—See *LMI*.

link state—The condition of an Open Shortest Path First (OSPF) link. See also *OSPF*.

Link-State Advertisement—See *LSA*.

link-state database—A database that contains Open Shortest Path First (OSPF) routing information. Link-state routing algorithms require that all routers within a domain maintain identical link-state databases, and that the databases describe the complete topology of the domain. An OSPF router's domain may be an Autonomous System (AS) or an area within an AS.

Based on the exchange of information among routers, OSPF routers create a link-state database, which is updated based on packet exchanges among the routers. Link-state databases are synchronized between pairs of adjacent routers. In addition, each OSPF router uses its link-state database to calculate a self-rooted tree of shortest paths to all destinations. The routing table is built from these calculated shortest-path trees. Externally derived routing data is advertised throughout the AS but is kept separate from the link-state data. Each external route can also be tagged by the advertising router, enabling the passing of additional information between routers on the boundary of the AS.

For example, suppose you have the network topology in Figure 25.

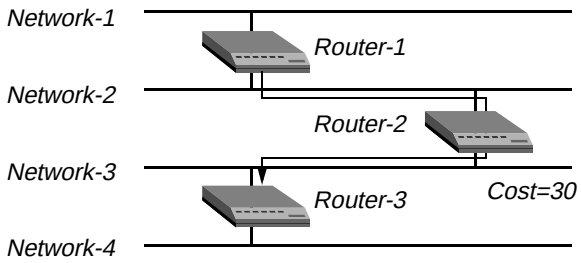

Figure 25. Sample OSPF network topology

The following table shows the information in the link-state databases of the three routers illustrated in Figure 25.

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

Each router builds a self-rooted shortest-path tree, and then calculates a routing table stating the shortest path to each destination in the AS (Figure 26, Figure 27, and Figure 28).

Figure 26. Shortest-path tree and resulting routing table for Router-1

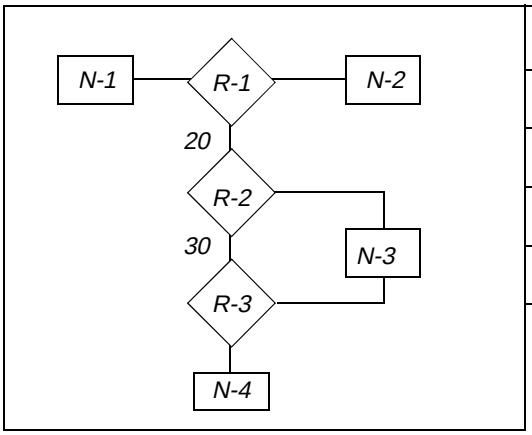

Destination	Next Hop	Metric
Network-1	Direct	0
Network-2	Direct	0
Network-3	Router-2	20
Network-4	Router-2	50

Figure 27. Shortest-path tree and resulting routing table for Router-2

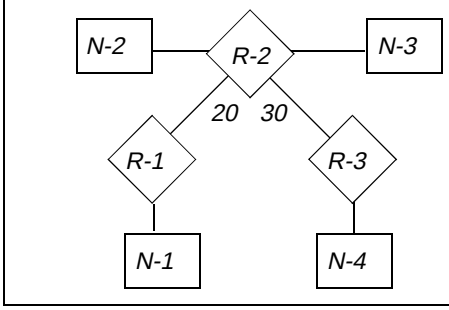

Destination	*Next Hop*	*Metric*
Network-1	Router-1	20
Network-2	Direct	0
Network-3	Direct	0
Network-4	Router-2	30

Figure 28. Shortest-path tree and resulting routing table for Router-3

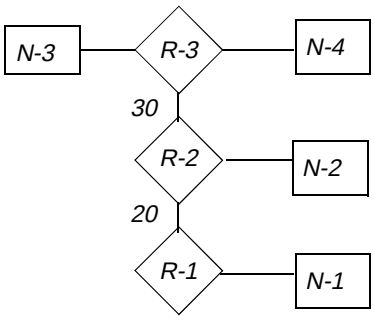

Destination	*Next Hop*	*Metric*
Network-1	Router-2	50
Network-2	Router-2	30
Network-3	Direct	0
Network-4	Direct	0

See also *adjacency*, *AS*, *OSPF*.

link-state metric—A metric that takes into account a variety of link conditions, such as the reliability or speed of the link, when determining the best path to a destination network. Open Shortest Path First (OSPF) is a link-state protocol. Another term for link-state metric is *interface cost*. Compare with *distance-vector metric*. See also *OSPF*.

Link-State-Request packet—An Open Shortest Path First (OSPF) request for an updated database. To make routing decisions, OSPF uses a link-state database of the network and propagates only changes to the database. See also *link-state database, OSPF*, *routing*.

Link-State-Update packet—See *LSU packet*.

List Attempt—See *DNS List Attempt*.

listening pattern—In an X.25/T3POS configuration, a called address that the Data Terminal Equipment (DTE) is expecting. Calls initiated by the host are answered by the DTE connecting to the T3POS PAD and listening for host-initiated calls. The host must send a called address matching the pattern the DTE is listening for. This pattern does not need to be a complete X.121 address, but could be a sub-pattern (including wildcard characters). See also *DTE, X.25/ T3POS, X.121*.

Livingston RADIUS—A version of RADIUS developed by Livingston Enterprises (now Lucent Technologies Remote Access Business Unit).

LLC—Logical Link Control. In the IEEE's Local Area Network/Reference Model, LLC denotes a sublayer above the Media Access Control (MAC) sublayer. Combined, the LLC and MAC sublayers are equivalent to the Data Link layer in the OSI Reference Model. They give higher-level protocols access to the physical media. See also *MAC, OSI Reference Model*.

LMI—Link Management Interface. LMI is a synchronous polling scheme used for the link management of a Frame Relay channel. The user polls the network to obtain status information about the Permanent Virtual Circuits (PVCs) configured on the channel. LMI exchanges this information by means of DLCI 1023. See also *DLCI, Frame Relay, PVC.*

LNS—L2TP Network Server. An LNS performs the following functions:

- Responds to requests by L2TP Access Concentrator (LAC) units for the creation of tunnels.
- Encapsulates and forwards all traffic from the private network to clients via the tunnel.
- De-encapsulates traffic received from an established tunnel, and forwards it to the private network.
- Disconnects tunnels on the basis of requests from the LAC.
- Disconnects tunnels on the basis of the expiration of the value you set. You can also manually disconnect tunnels from the LNS by means of SNMP, the terminal-server Kill command, or the DO Hangup command.

See also *L2TP, LAC.*

Local Area Network—See *LAN*.

Local Area Transport—See *LAT*.

local device—A device directly connected to the Ascend unit or residing on the local Ethernet.

local DNS table—Local Domain Name System table. A local DNS table resides in RAM, and contains up to eight hostnames and IP addresses. The MAX consults the local DNS table for address resolution only if requests to the DNS server fail. The local table acts as a safeguard to ensure that the MAX can resolve the local set of DNS names if all DNS servers become unreachable or go down. The table can contain up to 35 IP addresses per hostname entry. Following is a sample DNS table:

```
Name                          IP Address       # Reads Time of last read
____________________________ _________________ ________ _________________

 1: "barney"                  200.65.212.12 *   2       Feb 10 10:40:44 98
 2: "rafael"                  200.65.212.23     3       Feb 10  9:30:00 98
 3: "donatello"               200.65.212.67     1       Feb 11 11:41:33 98
 4: "wheelers"                200.65.212.9      1       Feb 12  8:35:22 98
 5: "tiktok"                  200.65.212.148    4       Feb 12  7:01:01 98
 6: ""                        -------           -       ---
 7: "wilma"                   200.65.212.8      10      Feb 15 10:02:58 98
 8: ""                        -------           -         ---
```

The table contains the following fields:

Field	Description
`Name`	Hostname.
`IP address`	IP address. An asterisk (*) indicates that the entry has been automatically updated by a DNS query.
`# Reads`	Number of accesses since the entry was created.
`Time of last read`	Time and date the entry was last accessed. The time and date appear only if Simple Network Time Protocol (SNTP) is in use. If SNTP is not in use, the field contains a row of hyphens.

See also *DNS*.

Local Exchange Carrier—See *LEC*.

local hunt-group number—A telephone number assigned to a hunt group for a single MAX. Compare with *global hunt-group number*. See also *hunt group*.

local loop—The connection between the telephone company's Central Office (CO) and its customers' home and business phones.

local loopback—A port diagnostic procedure in which data originating at the local site is looped back to its originating port without going out over the WAN. It is as though a data mirror were held up to the data at the WAN interface, and the data were reflected back to the originator. The AIM port on the MAX must be idle when you run the local loopback test. It can have no calls online. Compare with *remote loopback.* See also *analog loopback, digital loopback, loopback.*

local management—The process of managing a system when directly connected to the console, rather than over a WAN connection. Compare with *remote management*.

Local mode—A data-transfer mode for calls on an X.25/T3POS network. In Local mode, error recovery is performed locally. The MAX does not send supervisory frames (ACKs and NAKs) across the X.25 network. The T3POS PAD is responsible for sending supervisory frames to the T3POS Data Terminal Equipment (DTE). Compare with *Binary Local mode, Blind mode, Transparent mode*. See also *DTE, PAD, X.25/T3POS*.

local profile—A profile configured on the Ascend unit, in contrast to a user profile configured in RADIUS, TACACS, or TACACS+. See also *Connection profile, user profile*.

local user—A user at a device directly connected to the Ascend unit or residing on the local Ethernet.

logfile—A RADIUS file that contains error messages. You must create `logfile` yourself.

logical address—An address assigned by a network administrator to associate several devices with one another in a logical hierarchy or group. A router uses the logical address to help transmit a packet to its destination. An example of a logical address is an IP address. Compare with *hardware address*. See also *IP address, router*.

logical channel—A packet-switched communications circuit between two or more network hosts. Many logical channels can exist simultaneously on a single physical channel. See also *LCN*.

Logical Channel Number—See *LCN*.

logical interface—A Permanent Virtual Connection (PVC) endpoint on a Frame Relay network. The logical interface requires a Data Link Connection Indicator (DLCI). A DLCI uniquely identifies the logical endpoint of a Virtual Circuit (VC). See also *DLCI, Frame Relay network, PVC, VC*.

Logical Link Control—See *LLC*.

login prompt—The string used to prompt for a user name in the terminal-server interface when authentication is in use and an interactive user initiates a connection. See also *terminal server*.

login timeout—The number of seconds a terminal-server user can use for logging in. After the specified number of seconds, the login attempt times out. See also *terminal server*.

log level—The level of event information the MAX displays at the console.

Longitudinal Redundancy Check—See *LRC*.

loopback—A test in which a signal is sent to a destination and then returned to the sending device. See also *analog loopback, digital loopback, local loopback, remote loopback*.

loop start signaling—A type of signaling in which the Customer Premises Equipment (CPE) signals an off-hook condition by closing a relay at the Central Office (CO). Compare with *ground start signaling, wink-start signaling.*

LQM—Line Quality Monitoring. LQM is a feature that enables the MAX to monitor the quality of a link. When you enable LQM, the MAX counts the number of packets sent across the link and periodically asks the remote end how many packets it has received. Discrepancies are evidence of packet loss and indicate link-quality problems. The MAX can tear down and reestablish a call if the problems on the link exceed a specified threshold.

LRC—Longitudinal Redundancy Check. LRC is an error-detection method that adds one character, known as the Block Check Character (BCC), to the end of each packet. The system determines the value of the first bit of the BBC by counting the number of 1s in the first bits of all the characters in the packet, and then setting the first bit of the BBC to a 1 (one) if the sum is even, or to 0 (zero) if the sum is odd. The system determines the value of the second bit of the BBC by counting the number of 1s in the seconds bits of all the characters in the packet, and so forth. Compare with *CRC*.

LSA—Link-State Advertisement. An LSA is a packet that describes various aspects of an Open Shortest Path First (OSPF) route. Following are the available LSAs:

Type	Description
Type 1 (RTR)	Router-LSA that describes the collected states of the router's interfaces.
Type 2 (NET)	Network-LSA that describes the set of routers attached to the network.
Types 3 and 4 (STUB)	Summary-LSA that describes point-to-point routes to networks or Area Border Routers (ABRs).
Type 5 (ASE)	AS-external-LSA that describes routes to destinations external to the AS. An AS-external-LSA can also describe a default route for the AS.

See also *AS*, *ASE*, *ASE Type-5*, *external LSA*, *internal LSA*, *OSPF*, *point-to-point link*, *route*, *router*.

LSU packet—Link-State-Update packet. An LSU packet is exchanged between Open Shortest Path First (OSPF) routers for the purpose of updating link-state databases. See also *OSPF*, *router*.

LT mode—Line Termination mode. LT mode is the termination point of a WAN connection. Typically, it is the Customer Premises Equipment (CPE).

M

MAC—Media Access Control. In the IEEE's Local Area Network/Reference Model, MAC denotes a sublayer below the Logical Link Control (LLC) sublayer. Combined, the LLC and MAC sublayers are equivalent to the Data Link layer in the OSI Reference Model. They give higher-level protocols access to the physical media. See also *LLC*, *MAC address*, *OSI Reference Model*.

MAC address—The 6-byte hexadecimal address that the manufacturer assigns to the Ethernet controller for a port. The MAC address is also called an *Ethernet address*. See also *hardware address*, *MAC*.

Main Status Menu window—A window containing a hierarchical menu that displays an entry for each line or installed card in the MAX. The structure of the Main Status Menu exactly follows the Main Edit Menu (the top-level configuration menu). Following is a sample Main Status Menu window:

```
|--------------------|
|Main Status Menu    |
|>00-000 System     ^|
| 10-000 Net/T1      |
| 20-000 Net/T1     v|
|--------------------|
```

When the window that displays the Main Status Menu is active, the menu works like the Main Edit Menu. Use the arrow keys to scroll to a particular status menu. Then, press Return to open that menu and ESC to close it.

Management Information Base—See *MIB*.

manager—An application that receives Simple Network Management Protocol (SNMP) information from an agent. An agent and manager share a database of information, called the Management Information Base (MIB). An agent can use a message called a traps-PDU to send unsolicited information to the manager. A manager that uses the Ascend Enterprise MIB can query the MAX, set parameters, sound alarms when certain conditions appear, and perform other administrative tasks. See also *agent*, *community name*, *MIB*, *SNMP*, *traps-PDU*.

mark parity—A synonym for *odd parity.* See *parity*.

mark pattern—A pattern of 1s (11111111) that the MAX can use as the idle indicator on a dynamic AIM call. Compare with *flag pattern*.

mask—In a generic filter, a 12-byte value the MAX applies to a packet before comparing its contents to the value you indicate in a filter specification. The mask hides the bits that appear behind each binary 0 (zero). A mask of all ones (FF:FF:FF:FF:FF:FF:FF:FF) masks no bits, so the full specified value must match the packet contents. See also *generic filter*.

MAX—A multiprotocol WAN access router designed for central-site remote-access applications. The MAX offers the following features:

- Digital WAN access and services
- Digital and analog modems that dial in over channelized T1/PRI and E1/PRI access lines
- IP and IPX routing, bridging, and terminal-server functions
- Multiple-call aggregation for Bandwidth-on-Demand
- Multiple security methods
- Sophisticated management and control features
- Up to six optional slot cards

MAXDial—A software application that allows a LAN user without a modem to use the modems attached to the MAX and initiate fax and data transmissions. MAXDial enables users to make efficient use of the hardware and telephone lines already in operation.

Maximum Receive Unit—See *MRU*.

Maximum Reconstructed Receive Unit—See *MRRU*.

Maximum Transfer Unit—See *MTU*.

MBONE—Multicast Backbone. The MBONE is a virtual network layered on top of the Internet to support IP multicast routing across point-to-point links. Because multicasting is a fast and inexpensive way to communicate information to multiple hosts, the MBONE is used for transmitting audio and video on the Internet in real time.

The MBONE consists of groups of networks called *islands*. These islands are connected by tunnels and support IP. When the MAX accesses an MBONE network, it starts receiving MBONE multicasts. It resends the multicast packets to all of its own clients connected to it for MBONE service. The clients wanting MBONE service must implement Internet Group Membership Protocol (IGMP).

To the MBONE, the MAX looks like a multicast client. It responds as a client to IGMP packets it receives from an MBONE router. The MBONE router can reside on the MAX unit's Ethernet interface or across a WAN link. If the router resides across a WAN link, the MAX can respond to multicast clients on its Ethernet interface as well as across the WAN.

To multicast clients on a WAN or Ethernet interface, the MAX looks like a multicast router, although it simply forwards multicast packets on the basis of group memberships. See also *multicast, multicast forwarding, multicast heartbeat, multicast network, multicast rate limit, point-to-point link.*

MBONE interface—The location on the MAX that connects to an MBONE router. See also *MBONE router*.

MBONE router—A router that directs multicast packets to a group of clients on a subscription list. See also *MBONE, MBONE interface, multicast, multicast forwarding, multicast heartbeat, multicast network, multicast rate limit.*

MD5—A message-digest algorithm for security applications in which a large message is compressed and then signed with a private key. MD5 takes a message of an arbitrary length and creates a 128-bit message digest. See also *authenticator field, CHAP*.

menu mode—A mode in which the terminal server presents a banner message and a menu of hosts. In menu mode, a user cannot enter terminal-server commands, but can connect by means of Telnet, Rlogin, or raw TCP to the hosts you specify. The MAX authenticates the user's login name and password, and then displays a text-based menu such as the one shown in Figure 29.

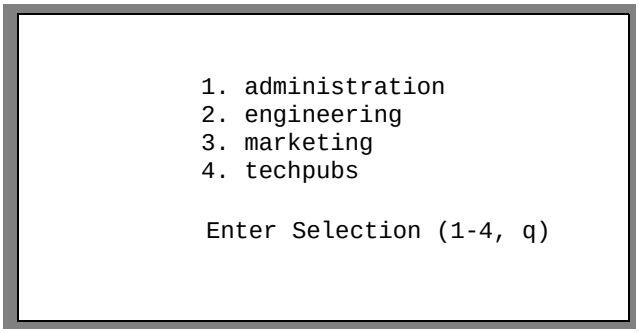

Figure 29. Menu mode

Users can Telnet to the specified host by pressing 1, 2, 3, or 4, or can quit the menu by pressing Q. Quitting the menu terminates the connection.

If you configure the menu locally, you can specify up to four hosts. If you configure the menu in RADIUS, you can configure up to ten. Compare with *command mode*, *immediate mode*.

Merit RADIUS—A version of RADIUS developed by Merit to design and implement an authentication, authorization and accounting system. It is built on the RADIUS protocol originally developed by Livingston Enterprises (now Lucent Technologies Remote Access Business Unit). See also *Livingston RADIUS*, *RADIUS*.

message—Data transmitted from one location to another with a header field, information field, and trailer. Often used interchangeably with *packet* and *frame*.

metric—A value that determines how quickly a packet can reach its destination. Routing Information Protocol (RIP) and the Open Shortest Path First (OSPF) protocol use different types of metrics.

- RIP is a distance-vector protocol, which uses a hop count to select the shortest route to a destination network. RIP always uses the lowest hop count, regardless of the speed or reliability of a link.
- OSPF is a link-state protocol. OSPF can take into account a variety of link conditions, such as the reliability or speed of the link, when determining the best path to a destination network.

See also *hop count*, *OSPF*, *preference*, *RIP*, *route*.

MHRP—Mobile Host Routing Protocol. MHRP is a protocol designed to support the mobility of a host. Using MHRP, a developer can design a product allowing continuous network connectivity for traveling computer users.

MIB—Management Information Base. A MIB is a Simple Network Management Protocol (SNMP) database of information available to network management programs. An agent creates a MIB. A network manager queries the MIB for information, and might create a MIB of its own. The MIB on the agent contains machine-specific information. The manager's MIB has more general information. The MAX supports a number of different MIBs. See also *agent, manager, SNMP*.

Microcom Networking Protocol—See *MNP*.

Microsoft CHAP—See *MS-CHAP*.

Microsoft Stac—The version of the Stac LZS compression method implemented by Windows 95. Compare with *Stac compression*, *Stac-9 compression.*

MNP—Microcom Networking Protocol. MNP is a communications hardware protocol developed by Microcom, Inc. Used by many high-speed modems, MNP supports several classes of communication. A modem can support more than one class.

Class 4 provides error detection, and can vary the modem's transmission speed in accordance with the quality of the line. Class 5 offers data compression, and can enable a device to double its transmission speed. Class 6 tries to detect the highest transmission speed supported by the modem at the other end of the connection, and then attempts to transmit data at that speed. Class 10-EC offers error detection and correction for cellular communication. The most commonly used MNP classes are Class 4 and Class 5, also called MNP-4 and MNP-5, respectively.

See also *compression*, *V.42*.

Mobile Client—A user or device that accesses a private Home Network across the Internet through an Ascend Tunnel Management Protocol (ATMP) tunnel. Using ATMP, a traveling salesperson or technical support specialist can dial into a local ISP and log into his or her Home Network. See also *ATMP*, *Home Network*.

Mobile Host Routing Protocol—See *MHRP*.

modem—MOdulator/DEModulator. A modem is Data Circuit-terminating Equipment (DCE) installed between Data Terminal Equipment (DTE) and an analog transmission channel, such as a telephone line. (A DTE refers to a device that an operator uses, such as a computer or a terminal. The DCE connects the DTE to a communications channel, such as a telephone line.) A modem takes digital data from a DTE, translates (or modulates) the 1s and 0s into analog form, and sends the data over the channel. The receiving modem demodulates the analog signal into digital data and sends it to the DTE to which it is attached. Compare with *digital modem.* See also *analog data*, *analog signal*, *DCE*, *digital data*, *digital signal*, *DTE*.

modem dialout—A feature that enables local users to connect to the terminal server by means of Telnet, and then issue AT commands to the digital modem as though connected locally to the modem's asynchronous port. You can configure the MAX for modem dialout to any Telnet port, or you can specify direct access to a particular Telnet port for immediate dialout service. When you specify direct access, you enable the *immediate modem* feature. See also *digital modem*, *immediate modem service*.

modem ringback—See *ringback tone*.

modem speed—The data rate for analog-modem transmissions. Most modem users currently receive data and voice across the local loop by means of a modulation process based on protocols such as V.32bis or V.34. These protocols limit users to connection speeds between 14.4 Kbps and 33.6 Kbps. Recently introduced 56K technology enables analog-modem users to download data at 56 Kbps over the local loop. See also *56K modem*, *modem*, *V.32bis*, *V.34*.

Mode Switch Frame—See *MSF*.

module—On the MAX, hardware that provides functionality for the base system or that plugs into an expansion slot. A *virtual module* reflects a function of the base system. Virtual module 0 manipulates overall system functions. Virtual module 1 is the Net/T1 module, which manipulates the base system's two-line T1 PRI network interface. Virtual module 2 is the Host/Dual module, which manipulates the base system's two AIM ports. A *real module* (3–8) plugs into an expansion slot in the MAX. See also *Host/Dual module*, *Net/T1 module*.

MP—Multilink Protocol. MP uses the encapsulation defined in RFC 1990, enabling the MAX to interact with MP-compliant equipment from other vendors. MP is an extension of Point-to-Point Protocol (PPP) and supports the ordering of data packets across multiple channels. The MP connection shown in Figure 30 has two channels.

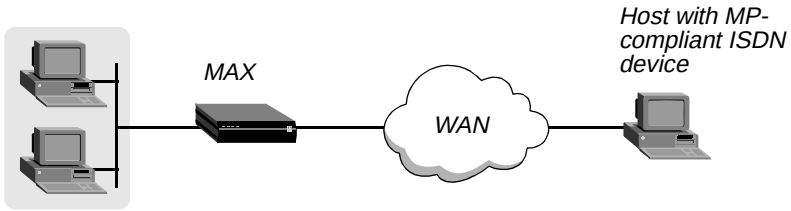

Figure 30. Multilink Protocol (MP) connection

If you configure an MP connection and the MAX cannot successfully negotiate the session, the unit falls back to using single-channel PPP. Compare with *MP+*, *PPP*.

MP+—Multilink Protocol Plus. MP+ uses Point-to-Point Protocol (PPP) encapsulation with Ascend extensions, as described in RFC 1934, to extend the capabilities of Multilink Protocol (MP). MP+ supports session and bandwidth management, enabling the MAX to connect to another Ascend unit by means of multiple channels. Using MP+, you can combine up to 30 individual channels into a single high-speed connection.

The connection in Figure 31 uses MP+ encapsulation between two MAX units.

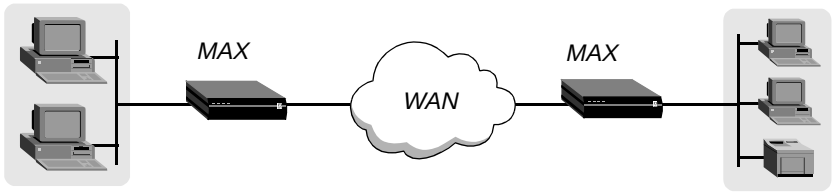

Figure 31. Multilink Protocol Plus (MP+) connection

MP+ consists of two components: a low-level channel-identification, error-monitoring, and error-recovery mechanism, and a session-management level for supporting bandwidth modifications and diagnostics. MP+ enables the MAX to add or remove channels from a connection as bandwidth needs change without disconnecting the link. This capability is called Dynamic Bandwidth Allocation (DBA).

Both the dialing side and the answering side of the link must support MP+. If you configure an MP+ connection and the MAX cannot successfully negotiate the link, the unit falls back to MP. If the MAX also fails to negotiate an MP connection, the unit falls back to using single-channel PPP.

MP+ calls cannot combine an ISDN BRI channel with a channel on a T1 PRI line. Compare with *MP, PPP*. See also *DBA, ISDN BRI line, T1 PRI line.*

MPP—Multilink Protocol Plus. This acronym has been superseded by MP+. See *MP+*.

MRRU—Maximum Reconstructed Receive Unit. MRRU is a packet field that indicates that the system implements the Multilink Protocol (MP). The MRRU field is two octets, and specifies the maximum number of octets in the Information fields of reassembled packets. A system must be able to receive the full 1500-octet Information field of any reassembled PPP packet, although it may attempt to negotiate a different value. See also *MP*.

MRU—Maximum Receive Unit. An MRU is the largest packet that a host on a link can receive. Compare with *MTU*.

MS-CHAP—Microsoft CHAP. MS-CHAP is a close derivative of Challenge Handshake Authentication Protocol (CHAP). However, CHAP is designed to authenticate WAN-aware secure software, and is not widely used to support remote workstations, where an insecure plain text login might be required. MS-CHAP addresses this issue, and also integrates the encryption and hashing algorithms used on Windows networks. Microsoft Windows NT and LAN Manager platforms implement MS-CHAP. Compare with *CHAP*.

MSF—Mode Switch Frame. After the call has been established for an X.25/T3POS connection, the Data Terminal Equipment (DTE) sends the host an MSF in order to inform the host of the mode of the call.

MTU—Maximum Transfer Unit. An MTU is the largest packet that can be transmitted over a particular medium. If a packet's size exceeds the MTU, the packet must be fragmented or segmented, and then reassembled at the receiving end. Compare with *MRU*.

Mu-Law—See *U-Law*.

Multiband inverse-multiplexing card—A MAX card that has two or six user-selectable RS-449, V.35, or X.21 serial host ports with inverse multiplexing and RS-366 capability, V.25bis, or control-lead signaling. See also *RS-449*, *RS-366*, *V.25bis*, *V.35*, *X.21*.

multicast—A transmission method in which one device communicates with destination hosts by means of a single transmission to all recipients of a subscriber list. The multicast destination addresses are 224.0.0.0 to 239.255.255.255. See also *MBONE, multicast forwarding, multicast heartbeat, multicast network, multicast rate limit.*

Multicast Backbone—See *MBONE*.

multicast default route—A route to the MBONE interface on the MAX. When the MAX acts as a multicast forwarder, and finds that there is no member in a particular group, it forwards multicast traffic for that group to the MBONE interface. See also *MBONE, multicast forwarding*.

multicast forwarding—A process by which the MAX forwards traffic it receives on one of its Ethernet or WAN interfaces from an Multicast Backbone (MBONE) router. Figure 32 shows a multicast router on a WAN interface with both local and WAN multicast clients.

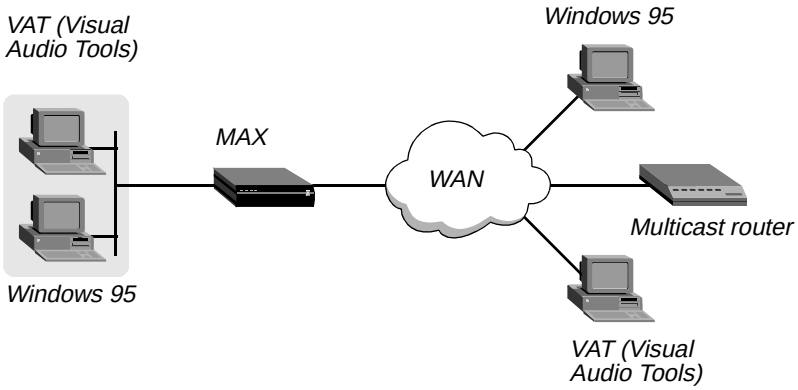

Figure 32. Forwarding multicast traffic on both Ethernet and WAN interfaces

To the MBONE, the MAX looks like a multicast client, and it responds as a client to Internet Group Membership Protocol (IGMP) packets it receives. The MAX resends the multicast packets to all of its own clients connected to it for MBONE service. The clients wanting MBONE service must implement IGMP.

Each Ethernet or WAN interface that supports multicasting must be configured to allow multicasting forwarding. When you do so, the MAX begins handling IGMP requests and responses on the interface. It does not begin forwarding multicast traffic until you set the multicast rate limit. See also *IGMP, MBONE, MBONE router, multicast, multicast heartbeat, multicast network, multicast rate limit.*

multicast group—A group of subscribers to whom a device sends multicast transmission. Membership in a multicast group is voluntary. Using Internet Group Membership Protocol (IGMP), you can configure an application on your PC to declare itself a member of a multicast group.

multicast heartbeat—A feature that enables you to monitor possible connectivity problems. Using the multicast heartbeat feature, you configure the MAX to poll continuously for multicast traffic. The MAX generates the following SNMP alarm trap if a traffic breakdown occurs:

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

Heartbeat monitoring is optional. It is not required for multicast forwarding. To set up heartbeat monitoring, you configure several parameters that define what packets will be monitored, how often the MAX polls for multicast packets, and what threshold must be reached for the MAX to generate an alarm. See also *MBONE, multicast, multicast forwarding, multicast network, multicast rate limit, SNMP*.

multicast network—A network in which a router sends packets to all addresses on a subscriber list. This type of network is different from both a unicast network (in which the router sends packets to one user at a time) and a broadcast network (in which the router sends packets to all users, whether they appear on subscription lists or not). The Multicast Backbone (MBONE) is an example of a multicast network. See also *MBONE, multicast, multicast forwarding, multicast heartbeat, multicast rate limit*.

multicast rate limit—A way to limit the rate at which the MAX accepts multicast packets from its clients. To begin forwarding multicast traffic on the MBONE interface, you must set the multicast rate limit to a number less than 100. For example if you set the limit to 5, the MAX accepts a packet from multicast clients on the interface every 5 seconds. Any subsequent packets received in that 5-second window are discarded. See also *MBONE, MBONE interface, multicast, multicast forwarding, multicast heartbeat, multicast network*.

multihomed host—A single Internet device connected to multiple data paths. Each link may reside on a different network.

Multilink Protocol—See *MP*.

Multilink Protocol Plus—See *MP+*.

multimode agent—A MAX that acts as either a Home Agent or a Foreign Agent on a tunnel-by-tunnel basis in an Ascend Tunnel Management Protocol (ATMP) configuration. In Figure 33, the MAX operates as a Home Agent for network B and as a Foreign Agent for network A.

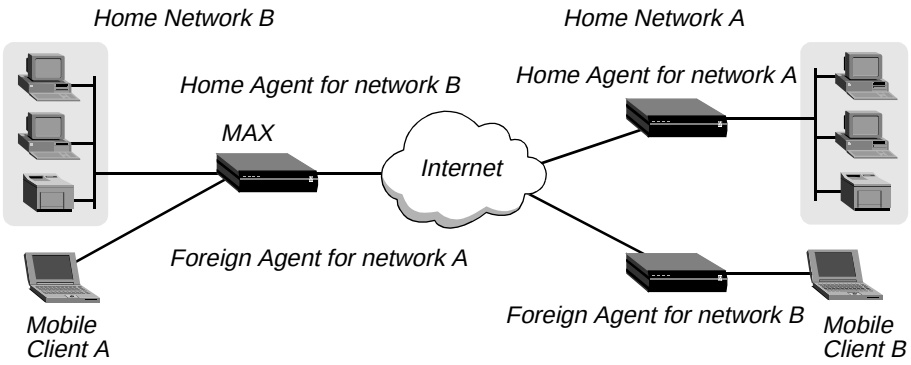

Figure 33. MAX *acting as both Home Agent and Foreign Agent*

See also *ATMP*, *Foreign Agent*, *Home Agent*.

multipath route—A static route that distributes the traffic load across multiple interfaces to a single destination. See also *route*, *static route*.

Multiple-address NAT—Multiple-address Network Address Translation. Multiple-address NAT provides a method of translating addresses for more than one host on the local network. The MAX borrows an official IP address for each host from a Dynamic Host Configuration Protocol (DHCP) server on the remote network (or on a network accessible from the remote network).

When you use multiple-address NAT, hosts on the remote network can connect to any of the official IP addresses that the MAX borrows from the DHCP server. If the local network must have more than one IP address visible to the remote network, you must use multiple-address NAT. If hosts on the remote network need to connect to specific hosts on the local network (not just specific services), you can configure the DHCP server to always assign the same address when that local host requests an address.

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

multiplexing—The technique of sending several signals on a communications medium at the same time.

multipoint link—A connection that links multiple hosts on a single line.

multipoint mode—A telephone service that provides a way for a single interface to have multiple telephone numbers.

multiprotocol routing—The ability to route multiple network protocols, including IP, IPX, and AppleTalk.

MultiRate—A data service on a circuit consisting of multiple B channels. The bandwidth of the circuit must be a multiple of 64 Kbps. For example, a user can dial a first call at 384 Kbps (using 6 B channels), and then dial a second call at 512 Kbps (using 8 B channels). MultiRate service is available over T1 PRI lines only. MultiRate is also known as the *Switched Nx64 data service*. See also *B channel, T1 PRI line*.

MultiVoice—An Ascend product that enables ordinary telephones to connect to other telephones using a public or private packet network. MultiVoice works by using a standard Voice-over-IP (VoIP) gateway that allows ordinary telephone calls to be transmitted across a packet network. The Ascend's VoIP gateway is known as the MultiVoice Gateway. An additional component, the MultiVoice Access Manager (also know as the *H.323 Gatekeeper*) provides phone-to-IP address translation. See also *MultiVoice Access Manager, MultiVoice Gateway, VoIP*.

MultiVoice Access Manager—A MultiVoice component that provides the following features:

- Phone-to-IP address translation
- Web-based administration interface
- PIN-based user authentication
- Voice Virtual Private Network (VPN) support
- Telephone number aliases
- Call Detail Reporting (CDR)
- Gateway and user database support
- Third-party billing system support

See also *MultiVoice, MultiVoice Gateway*.

MultiVoice Gateway—A MultiVoice component that supports the ITU-T H.323 standard for transmitting voice over an IP network. When a voice call is received at a local MultiVoice Gateway, the voice signal is packetized, compressed, and transmitted over the packet network using standard protocols and voice compression technologies. At the remote gateway, the process is reversed and the call is delivered over the remote packet network to its intended destination. See also *MultiVoice, MultiVoice Access Manager*.

N

nailed group—A group of nailed channels designated in an X.25 or Frame Relay profile.

nailed/MPP call—A Multilink Protocol Plus (MP+) call in which nailed-up channels can be augmented with switched channels if bandwidth is needed. The MAX must be the originator of the switched call.

A Nailed/MPP connection is established when its nailed-up or switched channels are connected end-to-end. The switched channels are dialed when the MAX receives an outbound packet for the remote end and cannot forward it across the nailed-up connection, either because those channels are down or because they are being fully used.

If both the nailed-up and switched channels in a Nailed/MPP connection are down, the connection does not reestablish itself until the nailed-up channels are brought back up or the switched channels are dialed. If a nailed-up channel fails, the MAX replaces that channel with a switched channel, even if the call is online with more than the minimum number of channels.

See also *MP+*, *nailed-up channel*, *switched channel*.

nailed-up channel—A channel on a line rented from the phone company for exclusive use, 24 hours per day, seven days per week. See also *nailed-up circuit*, *nailed-up line*.

nailed-up circuit—A permanent connection between endpoints over which two parties exchange data. The number of nailed-up channels must be the same at both ends of the connection. For example, if there are five nailed-up channels at the local end, there must be five nailed-up channels at the remote end. However, channel assignments do not have to match. For example, channel 1 may be switched at the local end and nailed up at the remote end. A nailed-up circuit is also known as a *private circuit* or a *leased circuit*. See also *nailed-up channel*, *nailed-up line*.

nailed-up line—A line rented from the phone company for exclusive use, 24 hours per day, seven days per week. The connection exists between two predetermined points and cannot be switched to other locations. A nailed-up line is also called a *leased line*. See also *nailed-up channel*, *nailed-up circuit*.

NAK—Negative Acknowledgment. A NAK is a packet sent from a receiver to a sender, informing the sender that data is missing or corrupt. When a device receives a packet, it sends back a packet to the sending device. If all the data arrived without corruption, the packet is an acknowledgment (ACK). If some of the data is missing or corrupt, a NAK results, and acts as a request that the sender retransmit the data. See also *ACK*.

name and password authentication—A form of authentication in which the MAX attempts to match a caller's user name and password to the parameters or attributes specified in a profile. If name and password authentication is required, the MAX first attempts to match the caller's name and password to a local Connection profile. If authentication succeeds using a local Connection profile, the MAX uses the parameters specified in the profile to build the connection.

If it cannot find a matching Connection profile, the MAX looks for a RADIUS, TACACS, or TACACS+ profile containing a matching name and password. If authentication succeeds using a RADIUS user profile, the MAX uses the specified RADIUS attributes to build the connection. The MAX can then forward the call to its bridge/router or other destination. If authentication succeeds using a TACACS or TACACS+ profile, the MAX must make a request to the server for information about the resources and services the user can access.

See also *authentication*, *Connection profile*, *RADIUS*, *RADIUS server*, *TACACS*, *TACACS+*, *user profile*.

Name Binding Protocol—See *NBP*.

NAS—Network Access Server. An NAS is a device that provides LAN and WAN access for network hosts. The MAX is an example of an NAS.

NAT for LAN—Network Address Translation for LAN. NAT for LAN is a feature that allows a Pipeline to connect a LAN to a remote network, even if devices on the LAN have addresses that are not valid for the remote network. The Pipeline translates between the local network addresses and the remote network addresses.

Access to public networks requires the use of an official IP address that is unique across the entire network. Typically, a central authority assigns a range of addresses, and a local administrator distributes them. If access to a public network is not necessary, the local manager can assign addresses as he or she sees fit, even if the addresses are unofficial or belong to another company.

Because the supply of addresses is rapidly diminishing, a company might not be able to get official addresses for its entire network. A site might already have unofficial addresses, but now needs access to the Internet, where an official address is required. For these reasons, you might need a facility for borrowing an official address and dynamically translating between the local and official addresses. NAT for LAN provides this facility.

Figure 34 illustrates a basic NAT for LAN setup.

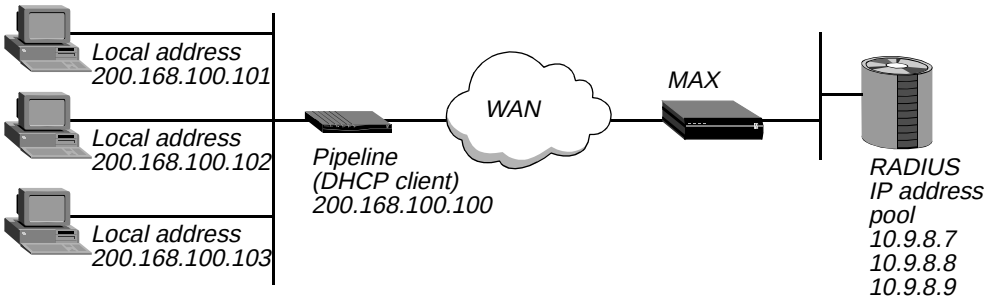

Figure 34. NAT for LAN setup

In Figure 34, the Pipeline itself does not have an address on the remote network. Therefore, clients can gain access to the Pipeline only from the local network, not from the WAN. When the first client on the LAN requests access to the remote network, the Pipeline gets an address for the client from the MAX. When subsequent clients request access to the remote network, the Pipeline sends the MAX a DHCP request packet, asking for an address. In return, the MAX sends an address from its IP address pool. The Pipeline uses the dynamic addresses it receives from the MAX to translate IP addresses on behalf of local clients.

As it receives packets on the LAN, the Pipeline determines whether the source IP address has a corresponding translated address. If so, the Pipeline translates the packet, and forwards it out the WAN. If the Pipeline has not assigned a translated address (and one is not pending), the Pipeline issues a new DHCP request for this IP address. While waiting for the MAX to offer an IP address, the Pipeline drops corresponding source packets. For packets it receives from the WAN, the Pipeline checks the destination address against its table of translated addresses. If the destination address exists and is active, the Pipeline forwards the packet. If the destination address does not exit, or is not active, the Pipeline drops the packet.

You can set up either multiple-address NAT or single-address NAT. See also *IP address*, *Multiple-address NAT*, *single-address NAT*.

National Center for Supercomputing Applications—See *NCSA*.

National ISDN-1—See *NI-1*.

National ISDN-2—See *NI-2*.

Navis—A service-management product for new public WANs. Navis delivers customer network-management support, scalability, service-provider operational-cost reduction, and end-to-end network control. Navis service management allows the Ascend hardware to be used to its full advantage, allowing for the creation of multiple services such as:

- Virtual Private Networks (VPNs) for dial-up and WAN capacity
- Varying classes of quality to meet varying traffic profiles and user business needs
- Bandwidth-on-demand to meet varying transport requirements
- Secure managed networks

There are three Navis software applications: NavisAccess, NavisCore, and NavisXtend. Each manages a specific set of Ascend hardware. See also *NavisAccess*, *NavisCore*, *NavisXtend*.

NavisAccess—An application that delivers superior management for the dial-up and dedicated portions of the network, providing extensive support for discovery and mapping, configuration, fault and performance management, and security. See also *Navis*, *NavisCore*, *NavisXtend*.

NavisCore—An application that operates in conjunction with HP OpenView to provide multiservice IP, Frame Relay, Asynchronous Transfer Mode (ATM) and Switched Multimegabit Data Service (SMDS) configuration and management of Ascend core switches from a single platform. See also *ATM*, *Frame Relay*, *IP*, *Navis*, *NavisAccess*, *NavisXtend*, *SMDS*.

NavisXtend—An application that extends the capabilities of NavisCore to bring additional features and cost efficiencies to switch network operations, such as provisioning automation, intelligent fault handling, and historical statistics storage. See also *Navis*, *NavisAccess*, *NavisCore*.

NBP—Name Binding Protocol. NBP is an AppleTalk protocol that enables you to make your application visible to users on an AppleTalk network. NBP associates the socket address assigned to a process or application to a name that contains three parts—the object, type, and zone fields. The object and type can be chosen by the developer, but the zone field specifies the zone in which the node resides. See also *AppleTalk*, *zone*.

NCP—NetWare Core Protocol. NCP is a protocol that allows an IPX server to respond to client requests. See also *IPX server*.

NCP—Network Control Protocol. NCP is a collection of protocols for setting up and configuring Network-layer protocols over PPP. See also *PPP*.

NCSA—National Center for Supercomputing Applications. The NCSA is the home of the first Web browser with a Graphical User Interface (GUI).

Near-End Block Error—See *NEBE*.

Nearest Server Query—See *IPX Nearest Server Query*.

NEBE—Near-End Block Error. A signal that the remote end sends to indicate that it has detected an error in the data it has transmitted. A block error is detected each time the calculated checksum of the data does not correspond to the control checksum transmitted in the successive superframe. One block error indicates that one superframe has not been transmitted correctly. No conclusion with respect to the number of bit errors can be drawn. Compare with *FEBE*.

negative acknowledgment—See *NAK*.

NetBIOS—Network Basic Input/Output System. NetBIOS is a protocol developed by IBM that provides network access to upper-layer programs. NetBIOS functionality includes the Session, Presentation, and Application layers of the OSI Reference Model, and provides naming services, connectionless best-effort datagram delivery, and support for Virtual Circuits (VCs). See also *OSI Reference Model*, *VC*.

Net/BRI port—On the MAX, a DTE port that provides a point-to-point ISDN BRI connection with another device. The MAX has eight Net/BRI ports. From the point of view of the MAX, pins 3 and 6 transmit on the Net/BRI interface, while pins 4 and 5 receive on the Net/BRI interface. See also *Host/BRI port*.

Net/E1 port—An E1 port that provides a point-to-point connection between the MAX and another device. The MAX has four Net/E1 or Net/T1 ports. See also *Net/T1 port*.

Net/T1 module—On the MAX, a module that manipulates the base system's two-line T1 PRI network interface. See also *Host/Dual module*, *module*.

Net/T1 port—A T1 port that provides a point-to-point connection between the MAX and another device. The MAX has four Net/E1 or Net/T1 ports. See also *Net/E1 port*.

NetWare Core Protocol—See *NCP*.

NetWare server—See *IPX server*.

network—A group of computers, often called *hosts*, *nodes*, or *stations*, that are connected to each other for the purpose of sharing files and other resources. Each computer has a Network Interface Card (NIC) that enables it to gain access to the network. Each host can have one or more peripherals (such as a fax modem or printer) attached to it. Each peripheral can be shared with other network users, or can remain private to the individual computer.

Network Access Server—See *NAS*.

network adapter—See *NIC*.

network address—An address shared by all the hosts on the same physical network.

Network Address Translation for LAN—See *NAT for LAN*.

network alignment—A method of setting up IP address pools for pool summary. When you perform network alignment, you make sure that the first address in the pool is the first host address, and that the maximum number of entries you specify is two fewer than the total number of addresses in the pool. See also *IP address*, *pool summary*.

Network Basic Input/Output System—See *NetBIOS*.

network board—See *NIC*.

Network Control Protocol—See *NCP*.

Network File System—See *NFS*.

Network Information Center—See *InterNIC*.

Network Information Service—See *NIS*.

Network Interface Card—See *NIC*.

Network layer—A layer in the OSI Reference Model. The Network layer provides address resolution and routing protocols. Address resolution enables the Network layer to determine a unique network address for a node. Routing protocols allow data to flow between networks and reach their proper destination. Examples of Network-layer protocols are Address Resolution Protocol (ARP), Datagram Delivery Protocol (DDP), Internet Control Message Protocol (ICMP), Interior Gateway Protocol (IGP), Internet Protocol (IP), Internetwork Packet Exchange (IPX), and Packet Layer Protocol (PLP). See also *ARP*, *DDP*, *ICMP*, *IGP*, *IP*, *IPX*, *OSI Reference Model*, *routing*.

network number—See *IP network number*, *IPX network number*.

Network News Transfer Protocol—See *NNTP*.

network port—A T1 or E1 channel. The MAX always places or receives calls on a network port.

network range—A contiguous range of integers (from 1 to 65, 199) assigned to an AppleTalk network. Each network range must be unique. No two networks can use the same range, and no two network ranges can overlap. In order for the MAX to receive calls from dial-in AppleTalk Remote Access (ARA) clients, you must define a virtual AppleTalk network by specifying a network range. Each number in the range can be associated with up to 253 nodes, so the range determines how many AppleTalk clients can dial into the MAX. For example, a network with a range from 1000 to 1002 could support up to 2 x 253, or 506 clients. See also *AppleTalk routing*, *ARA*, *virtual AppleTalk network*.

Network Service Provider—See *NSP*.

network switch—A network device that selects a path or circuit for sending data to its next destination.

network tone cut-through—A feature that provides answer supervision support for MAX gateways that use non-PRI trunks. Network tone cut-through enables each MAX to pass call-progress tones across the IP network for Voice-over-IP (VoIP) calls. Call-progress tones generated by a distant Public Switched Telephone Network (PSTN) are passed between the two MAX gateways processing the call. When a Far End Gateway receives call progress tones from the PSTN, the tones are stored as voice frames, then transmitted across the IP network in Real-Time Transport Protocol (RTP) packets. Upon receiving the RTP packets, the Near End Gateway decodes and sends these tones to the calling endpoint. See also *PSTN*, *RTP*, *VoIP*.

Network-to-Network Interface—See *NNI*.

Network User Identification—See *NUI*.

Network Virtual Terminal—See *NVT*.

Network Virtual Terminal ASCII—See *NVT ASCII*.

Network Voice Protocol—See *NVP*.

next-hop router—The router that is one hop away from another device. See also *hop*, *router*.

NFAS—Non-Facility Associated Signaling. NFAS is a special case of ISDN signaling in which two or more T1 PRI lines use the same D channel, and you can add a backup D channel. It is required for the Switched-1536 data service. Because all 24 channels of the T1 PRI line carry user data, the D channel must be on another line. See also *D channel*, *Switched-1536*, *T1 PRI line*.

NFS—Network File System. NFS is an Application-layer protocol, developed by Sun Microsystems, for sharing and transferring remote files on UNIX or other types of networks. See also *Application layer*.

NI-1—National ISDN-1. NI-1 is a standard service type for an ISDN phone line.It was created so that users would not have to know what kind of switch they were connected to in order to buy compatible equipment. All the Regional Bell Operating Companies (RBOCs) support NI-1, though a new standard, National ISDN-2 (NI-2), was recently adopted.

NI-2—National ISDN-2. A new switch type that represents a more comprehensive standard than National ISDN-1 (NI-1). See also *NI-1*.

NIC—See *InterNIC*.

NIC—Network Interface Card. A NIC enables a PC to connect to a network. The NIC uses drivers to communicate with the host's networking software, and interacts with the physical media that connects the host to other computers. A NIC is also called a *LAN adapter, network adapter*, or *network board*.

NIS—Network Information Service. Along with the Network File System (NFS), the NIS is a method of creating a distributed database system in order to centralize common configuration files, such as the UNIX password file (`/etc/passwd`) and the hosts file (`/etc/hosts`). An NIS server manages copies of the database files, and NIS clients request information from them. NIS was developed by Sun Microsystems. See also *NFS*.

NNI—Network-to-Network Interface. NNI is a standard that defines the interface between two Frame Relay switches located in a private or public network. NNI operation enables the MAX to act as a Frame Relay switch communicating with another Frame Relay switch.

The MAX uses NNI procedures to inform its peer switch about the status of Permanent Virtual Circuit (PVC) segments from its side of the Frame Relay network, and to communicate information about the integrity of the datalink between them. The procedure is bidirectional. The switches act as both the user side (DTE) and network side(DCE) in that they both send Status Enquiries and respond to them.

Figure 35 shows an example of the MAX with NNI interfaces.

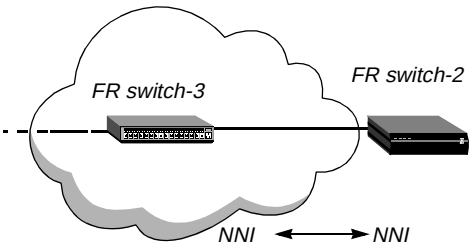

Figure 35. Frame Relay NNI interface

Compare with *UNI*. See also *DCE*, *DTE*, *Frame Relay*, *Frame Relay network*, *Frame Relay switch*, *PVC*.

NNTP—Network News Transfer Protocol. NNTP is the most commonly used protocol for exchanging news on Usenet newsgroups.

node—See *host*.

node number—A value assigned to a host on a network. The node number can be hardcoded in the Network Interface Card (NIC), or assigned by means of jumper settings. It is unique among all the hosts on a local, physical network. The address for a host also contains the network address shared by all the hosts on the local network. See also *host*, *network address*, *NIC*.

noise—On a communications medium, electrical or magnetic interference that can degrade the quality of a signal.

nonextended AppleTalk network—An AppleTalk network that contains a maximum of 254 nodes and is assigned a single network number. Compare with *extended AppleTalk network*. See also *AppleTalk*.

Non-Facility Associated Signaling—See *NFAS*.

nonseed router—An IPX or AppleTalk router that acquires its network configuration from another router on the network. Compare with *seed router*. See also *AppleTalk routing*, *IPX router*.

Nonvolatile Random Access Memory—See *NVRAM*.

normal area—An Open Shortest Path First (OSPF) area that allows Type-5 Link-State Advertisements (LSAs) to be flooded throughout it. Area Border Routers (ABRs) advertise external routes as Type-5 LSAs. A normal area is the default for the MAX. If you change the default for one interface on the unit, you must change it for all interfaces, because the MAX does not currently perform ABR functions. Compare with *NSSA*, *stub area*. See also *ABR*, *area*, *ASE Type-5*, *external route*, *LSA*, *OSPF*, *router*, *routing*.

Northern Telecommunications, Inc.—See *NTI*.

Not So Stubby Area—See *NSSA*.

NSP—Network Service Provider. An NSP is a company that provides Internet connectivity to Internet Service Providers (ISPs) and other organization requiring high-speed access to the Internet. See also *ISP*.

NSSA—Not So Stubby Area. An NSSA is an Open Shortest Path First (OSPF) area that does not receive or originate Type-5 Link-State Advertisements (LSAs), and that imports Autonomous System (AS) external routes in a limited fashion. OSPF version 2 defines a new Type-7 LSA for NSSAs.

For NSSAs, all routes imported to OSPF have the P-bit set (P stands for *propagate*). When the P-bit is enabled, Area Border Routers (ABRs) translate Type-7 LSAs to Type-5 LSAs, which can then be flooded to the backbone. These external routes are considered Type-7 LSAs.

Compare with *normal area*, *stub area*. See also *area*, *AS*, *ASE Type-5*, *ASE Type-7*, *external route*, *LSA*, *OSPF*.

NT1—Network Terminator Type 1. An NT1 is a terminating device for an ISDN BRI line. Installed at the subscriber's location, an NT1 provides line maintenance, timing, and echo cancellation. An NT1 can be a standalone device, or it can be built into other types of equipment. See also *ISDN BRI line*.

NTI—Northern Telecommunications, Inc. NTI is a company that manufactures a wide variety of telecommunications products.

NUI—Network User Identification. NUI is name/password combination that gives you access to a commercial packet-switched network.

numbered interface—In interface-based IP routing, a unique address assigned to one side of a connection. Assignment of a unique address is a requirement for some applications, such as Simple Network Management Protocol (SNMP). Figure 36 shows a local interface with two addresses, one of which is used for a numbered interface connection.

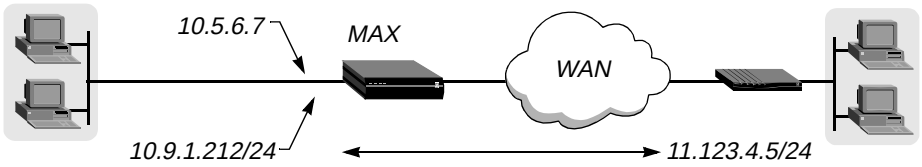

Figure 36. How numbered interfaces work

Reasons for using numbered interfaces include troubleshooting nailed-up point-to-point connections and forcing routing decisions between two links going to the same final destination. More generally, interface-based routing allows the MAX to operate more as a multihomed Internet host behaves. Compare with *system-based routing*, *unnumbered interface*. See also *interface-based routing*, *IP address*, *multihomed host*, *point-to-point link*, *SNMP*.

NVP—Network Voice Protocol. NVP is a protocol developed to enable the communication of real-time interactive voice over disparate computer networks.

NVRAM—Nonvolatile Random Access Memory. NVRAM is a type of memory that maintains its data contents across resets and power cycles. It is useful for storing configuration information across sessions. Data is written and erased in blocks, rather than byte-by-byte.

The MAX unit's system configuration is stored in the onboard NVRAM. Some error conditions may require that you clear the MAX configuration and reboot. When you clear NVRAM, the system is re-initialized and comes up unconfigured, just as it was when you first installed it. You can then restore the configuration from a recent backup.

NVRAM is also called *flash memory.* Compare with *DRAM*, *EEPROM*, *RAM*.

NVT—Network Virtual Terminal. An NVT is a bidirectional character device with a printer and a keyboard. The printer responds to incoming data, and the keyboard produces outgoing data sent over a Telnet connection. The code set is seven-bit ASCII in an eight-bit field. See also *NVT ASCII*, *Telnet session*.

NVT ASCII—The ASCII character set used with a Network Virtual Terminal (NVT). See also *ASCII*, *NVT*.

O

Octet—Eight data bits, also called a *byte*.

odd parity—See *parity*.

off hook—A state that results when you lift a telephone receiver, producing a busy signal.

offset—In a generic filter, the bytes from the start of a frame to the data in the packet to be tested against the filter. See also *generic filter*.

Open Shortest Path First—See *OSPF*.

Open Systems Interconnection Reference Model—See *OSI Reference Model*.

OSI Reference Model—Open Systems Interconnection Reference Model. The OSI Reference Model describes the layers of a network, details the functions of each layer, and explains how to connect communications devices on a LAN or WAN. Each layer provides services for the layer above it, and uses the services of the layer below it. From top to bottom, the seven layers are:

Layer	Description
Application	Provides applications with access to the network. File transfer, email, and network management software are examples of Application-layer programs. Protocols such as Simple Network Management Protocol (SNMP), Telnet, Rlogin, File Transfer Protocol (FTP), and File Transfer, Access, and Management (FTAM) provide Application-layer services.
Presentation	Responsible for presenting information in a format understandable to users and their applications. Data conversion, special graphics, compression, and encryption are some of the functions implemented at the Presentation layer.
Session	Synchronizes the data in a network connection, maintains the link until the transmission is complete, handles security, and makes sure that the data arrives in the proper sequence. Gateway communications are implemented at the Session layer. Examples of Session-layer protocols are AppleTalk Data Stream Protocol (ADSP), NetBEUI (an extension of NetBIOS), NetBIOS, and Printer Access Protocol (PAP).
Transport	Provides data transfer at the proper speed, quality, and error rate, ensuring reliable delivery. Examples of Transport-layer protocols are Transmission Control Protocol (TCP), User Datagram Protocol (UDP), and Sequenced Packet Exchange (SPX).

Layer	Description
Network	Provides address resolution and routing protocols. Address resolution enables the Network layer to determine a unique network address for a node. Routing protocols allow data to flow between networks and reach their proper destination. Examples of Network-layer protocols are Address Resolution Protocol (ARP), Datagram Delivery Protocol (DDP), Internet Control Message Protocol (ICMP), Interior Gateway Protocol (IGP), Internetwork Packet Exchange (IPX), Internet Protocol (IP), and Packet Layer Protocol (PLP).
Data Link	Creates, sends, and receives data packets appropriate for the type of network in use. Data Link-layer protocols include High-Level Data Link Control (HDLC), Link Access Procedure, Balanced (LAPB), Link Access Procedure, D channel (LAPD), Point-to-Point Protocol (PPP), and Serial Line Internet Protocol (SLIP).
Physical	Defines the electrical properties of the physical medium, and converts the data into a series of 0s and 1s for digital transmission. Examples of Physical-layer specifications include RS-232, RS-422, RS-423, RS-449, IEEE 802.3, and IEEE 802.5.

OSPF—Open Shortest Path First. OSPF is the next generation Internet routing protocol. The *Open* in its name refers to the fact that OSPF was developed in the public domain as an open specification. The *Shortest Path First* portion refers to an algorithm developed by Dijkstra in 1978 for building a self-rooted shortest-path tree from which routing tables can be derived. As a link-state protocol, OSPF an take into account a variety of link conditions, such as the reliability or speed of the link, when determining the best path to a destination network. OSPF uses a link-state database of the network and propagates only changes to the database. See also *link-state database*, *route*, *router*, *routing*.

outgoing call—A call that the MAX places to another device.

out-of-band management—A management method that uses a separate channel for diagnostic and administrative purposes (rather than a portion of each data channel).

out-of-band signaling—See *ISDN D-channel signaling*.

out-of-frame condition—A condition in which the T1 line cannot receive or transmit data because the MAX has lost the frame alignment on the received signal. See also *T1 line*.

output filter—A filter applied to an outgoing packet. See also *filter*, *packet filter*.

P

packet—A block of information containing a header, data, and trailer. Packets created at one level of the OSI Reference Model are inserted into lower-level packets. The format of a packet depends upon the protocol that creates it. A packet can be transmitted over a network or phone line. Compare with *frame*. See also *OSI Reference Model*, *packet field*.

Packet Assembler/Disassembler—See *PAD*.

packet field—A portion of a packet that contains a specific kind of information. For example, the data field in a packet contains the data being transmitted between applications. The header field can contain information identifying the packet type and any error-checking mechanisms. See also *packet*.

packet filter—A series of rules stating how the MAX is to handle certain types of packets. Each rule specifies a condition and an action to be taken if the condition is met. The MAX compares data in the packet to each condition, one condition at a time, until it finds a match between the data and one of the conditions. It then forwards or drops the packet, depending on the action specified for the condition.

When no filter is in use, the MAX forwards all packets. But when you apply a filter to an interface, you reverse that default. For security purposes, the unit no longer forwards non-matching packets automatically. It requires a rule that explicitly allows those packets to pass.

You can apply a packet filter to incoming packets, outgoing packets, or both. In addition, you can specify that the MAX forward or drop those packets that match the rules, or all packets *except* those that match the rules.

The MAX supports three types of packet filters: generic, IP, and IPX. You can apply a generic, IP, or IPX filter as either a data filter or a call filter. The MAX applies a data filter before a call filter.

See also *call filter*, *data filter*, *generic filter*, *IP filter*.

Packet Layer Protocol—See *PLP*.

packet-level inverse multiplexing—A method of inverse multiplexing in which the inverse multiplexer performs its function at the packet level by means of the Multilink Protocol (MP) or Multilink Protocol Plus (MP+). One data packet goes over the first circuit, the next goes over the second circuit, and so on, until all the data packets are distributed over all the available circuits. The receiving end adjusts for network-induced delay and reassembles the data packets into their proper order. This inverse multiplexing technique is also referred to as *load balancing*. Telecommuting applications use packet-level inverse multiplexing. Compare with *circuit-level inverse multiplexing*. See also *inverse multiplexer*, *inverse multiplexing*.

packet-radio communication—A technology for wireless modem access in which a system breaks a transmission into small packets that include a source address, destination address, and error-correction information. The packets are uplinked to a satellite, and then broadcast. The destination device receives only packets addressed to it. See also *error correction*, *wireless technology*.

Packet-Switched Data Network—See *PSDN*.

Packet-Switched Public Data Network—See *PSPDN*.

packet switching—A mode of data transfer in which packets are transmitted from a specific source to a specific destination using any available circuit. Packets may take different paths at the same time, and may not arrive in the order in which they were sent. Compare with *circuit switching*.

PAD—Packet Assembler/Disassembler. A PAD is an asynchronous terminal concentrator that enables several terminals (or other asynchronous devices) to share a single network line. See also *X.25/PAD*.

Palmtop port—A MAX port that connects to a handheld Palmtop control terminal. The Palmtop port provides access to the MAX unit's menu-driven interface. It runs at 9600 bps, eight bits per character, no parity, no flow control, one stop bit.

PAP—Password Authentication Protocol. PAP uses a two-way handshake method of establishing a caller's identity. Used only during the initial establishment of the data link, PAP is not a strong authentication method. Passwords travel across the line as plain text, so they are subject to eavesdroppers using software that monitors network information. Use PAP authentication only when the dial-in device does not support a stronger authentication method, such as Challenge Handshake Authentication Protocol (CHAP), or when the remote device requires a plain text password.

An extension of PAP adds the U.S. Data Encryption Standard (DES) cipher to data transmissions. The caller applies the encryption algorithm to a PPP packet and places the resulting cipher text in the information field of another PPP packet. The receiving end applies the inverse algorithm and interprets the resulting plain text as if it were a PPP packet that had arrived directly on the interface.

Compare with *CHAP*. See also *authentication*, *DES*, *password*, *PPP*.

PAP-Token authentication—An extension of Password Authentication Protocol (PAP) authentication. In PAP-Token authentication, the user authenticates his or her identity by entering a password (called a *token*). The token is derived from a hardware device, such as a hand-held token card. The MAX prompts the user for the token, possibly along with a challenge key. The MAX obtains the challenge key from a token-card server that it accesses through RADIUS. The token travels in the clear, but because it is a one-time-only password, the security risk is usually not serious. To authenticate the base channel of the connection, the token-card server uses the token that the user sends in response to the challenge.

PAP-Token is appropriate for single-channel, dial-out calls. It is not practical for multichannel calls, because any time that bandwidth requirements cause another channel to come up, the MAX challenges the user for another token.

Compare with *PAP*, *PAP-Token-CHAP authentication*. See also *RADIUS*, *token*, *token card*, *token-card authentication*, *token-card server*.

PAP-Token-CHAP authentication—An authentication method that uses PAP-Token to authenticate the base channel of Multilink Protocol Plus (MP+) call, and then a Challenge Handshake Authentication Protocol (CHAP) password to authenticate additional channels. The advantage of a PAP-Token-CHAP call over a PAP-Token call is that you need to verify only the initial connection by means of a hand-held token card. In a PAP-Token-CHAP call, the MAX uses CHAP to verify any additional channels. Compare with *CHAP*, *PAP*, *PAP-Token authentication*. See also *MP+*, *token*, *token card*, *token-card authentication*, *token-card server*.

parallel dialing—A way for the MAX to add channels to an outgoing call in multiples, rather than one at a time.

parity—In 7-bit communication, a way for a device to determine whether it has received data exactly as the sending device transmitted it. Each device must determine whether it will use even parity, odd parity, or no parity.

The sending device adds the 1s in each string it sends and determines whether the sum is even or odd. Then, it adds an extra bit, called a *parity bit*, to the string. If even parity is in use, the parity bit makes the sum of the bits even. If odd parity is in use, the parity bit makes the sum of the bits odd. For example, if a device sends the binary number 1010101 under even parity, it adds a 0 (zero) to the end of the byte, because the sum of the 1s is already even. However, if it sends the same number under odd parity, it adds a 1 to the end of the byte in order to make the sum of the 1s an odd number. (A synonym for even parity is *space parity*; a synonym for odd parity is *mark parity.*)

The receiving device checks whether the sum of 1s in a character is even or odd. If the device is using even parity, the sum of 1s in a character should be even. If the device is using odd parity, the sums of bits in a character should be odd. If the sum of the bits does not equal the parity setting, the receiving device knows that an error has occurred during the transmission of the data.

For special ASCII characters (128-256), eight bits are necessary to represent the data. In 8-bit communication, no parity bit is used. See also *ASCII*.

parity bit—An extra bit added to a string in 7-bit communication. The sending device adds the 1s in each string it sends and determines whether the sum is even or odd. Then, it adds the parity bit to the string. If even parity is in use, the parity bit makes the sum of the bits even. If odd parity is in use, the parity bit makes the sum of the bits odd.

password—A text string that a user must enter during the login process. Entering the proper password identifies the user as a person authorized to access network resources. Compare with *token*.

password prompt—The prompt that the terminal server displays when asking the user for his or her password. See also *password*, *terminal server*.

PBX—Private Branch Exchange. A PBX is an internal telephone network, such as those used in large offices, in which one incoming number directs calls to various extensions and from one office to another. See also *PRI-to-T1 conversion*.

PCM—Pulse Coded Modulation. PCM is a sampling technique for encoding a digital stream so that it contains a digitized version of the analog waveform sent by a device attached to a modem. The MAX can also convert outgoing data into analog waveforms, change these waveforms into a PCM-encoded digital stream, and send them to the network over a digital line. The network presents the data to the receiving modem in analog form over an analog line. The data looks exactly as it would appear if it had been sent by an analog-based modem.

There are two standards for coding the sample level. The U-Law standard is common in North America and Japan. Elsewhere, the A-Law standard is typically in use.

See also *analog line*, *digital line*, *modem*.

PCMCIA—Personal Computer Memory Card International Association. PCMCIA is a standard that supports the devices on a credit-card-sized board. The 1990 PCMCIA version 1.0 specification supports Type I cards for RAM, ROM, or NVRAM. The 1991 PCMCIA version 2.01 specification supports Type II cards for network and fax/modem functionality, and Type III cards. A Type III card provides a miniature hard drive for wireless networks. See also *NVRAM*, *RAM*, *ROM*, *wireless technology*.

PCMCIA card code—Code written to make use of PCMCIA-card functionality. See also *PCMCIA*, *PCMCIA flash card*, *PCMCIA interface*.

PCMCIA flash card—On the MAX, a standard card that extends existing flash memory. See also *PCMCIA*, *PCMCIA card code*, *PCMCIA interface*

PCMCIA interface—On the MAX, an interface that accepts a plug-in PCMCIA flash card. See also *PCMCIA*, *PCMCIA card code*, *PCMCIA flash card*.

PCS—Personal Communications Services. PCS is a wireless telephone service for mobile users, similar to cellular communication. It is also referred to as *digital cellular.* See also *cellular communication*, *wireless technology*.

PCS 1900—See *GSM 1900*.

PDU—Protocol Data Unit. A PDU is a packet created at any one of the OSI layers. See also *OSI Reference Model*.

peripheral—A device attached to a network, server, or workstation. Peripherals include CD-ROM drives, fax machines, hard drives, modems, optical drives, printers, and tape drives.

Personal Computer Memory Card International Association—See *PCMCIA*.

Personal Handyphone System—See *PHS*.

Personal Internet Access Forum Standard—See *PIAFS*.

per-user accounting—A way for network reseller to direct accounting information about specific users to a RADIUS server belonging to a particular ISP. A network reseller can serve many different ISPs, each with a different access policy. The reseller carries traffic for individual users, and must bill for usage according to the policies of the appropriate ISP. Per-user accounting facilitates this process. See also *accounting*, *ISP*, *RADIUS*.

per-user default route—The default route for IP packets coming from a particular user. The MAX uses the per-user default under either of the following circumstances:

- The next-hop address in the MAX unit's routing table is the default route for the system (destination 0.0.0.0).
- The normal routing logic fails to find a route and there is no system-wide default route.

The direct route can take place by means of a profile or an Ethernet connection. If the Ascend unit does not have a direct route, it drops the packets on the connection. The default value is 0.0.0.0. If you accept this value, the Ascend unit routes packets as the routing table specifies, using the system-wide default route if it cannot find a more specific route.

The per-user default route applies to all packets the MAX receives for a given profile, regardless of the specific IP source address. Therefore, you can use this feature when the profile belongs to another router, and all hosts behind that router use the default gateway. The MAX handles packets from other users or from the Ethernet network in the usual fashion. The global routing table is not altered. Therefore, when you diagnose routing problems with a profile that implements a per-user default route, an error in a per-user gateway address is not apparent from inspection of the global routing table.

See also *default route*, *hop*, *IP address*, *IP route*, *IP routing*.

PHS—Personal Handyphone System. PHS is a mobile phone system available in Japan. In addition to providing voice service, PHS allows data communication at rates of up to 32 Kbps, and can be used for Internet access. See also *PIAFS*.

Physical layer—The lowest layer in the OSI Reference Model. The Physical layer defines the electrical properties of the physical medium, and converts the data into a series of 0s and 1s for digital transmission. Examples of Physical-layer specifications include RS-232, RS-422, RS-423, RS-449, IEEE 802.3, and IEEE 802.5. See also *802.3*, *802.5*, *OSI Reference Model*, *RS-232*, *RS-422*, *RS-423*, *RS-449*.

PIAFS—Personal Internet Access Forum Standard. PIAFS is a protocol that handles connection negotiation, data transfer, and error correction for the Personal Handyphone System (PHS). See also *PHS*.

PID—Protocol Identifier. The PID is a group of bytes in the SNAP header of an IPX frame. The PID identifies the protocol in use for the transmission. See also *IPX frame*, *SNAP*.

PIFSA-16 card—A MAX card that provides up to eight PIFSA WAN sessions. You can install a maximum of six PIFSA-16 cards in the MAX.

Ping—A command that sends an Echo request in order to test whether a remote network device is accessible. If the remote device is properly connected, it receives the request and sends back an Echo Reply. Certain version of the Ping command can also determine the amount of time necessary to receive the Echo Reply, and the number of replies lost in transmission. See also *Echo*.

Plain Old Telephone Service—See *POTS*.

PLP—Packet Layer Protocol. A protocol that specifies the rules for full-duplex data transmission between a sending device and a receiver on an X.25 network. See also *full duplex*, *X.25*.

Plug and Play—See *PNP*.

PNP—Plug and Play. PNP describes the process of plugging a device into a computer and having the computer recognize it without user configuration.

PNS—PPTP Network Server. The device acting as the endpoint of a PPTP tunnel. See also *PPTP*.

Point of Presence—See *POP*.

point-to-point link—A connection that does not make any use of intervening devices. A point-to-point link can connect two hosts on the same network, or two networks across the WAN.

Point-to-Point Protocol—See *PPP*.

Point-to-Point Tunneling Protocol—See *PPTP*.

poison—Denotes the MAX unit's ability to stop advertising ("poison") IP dialout routes if it temporarily loses the ability to dial out. See also *IP route*.

poison reverse—A policy for handling RIP update packets that include routes received on the same interface on which the update was sent. Using a poison-reverse policy, the MAX propagates routes back to the subnet from which they were received and assigns them a metric of 16. See also *metric*, *RIP*, *subnet*.

pool summary—A configuration in which the router advertises a single route for the network you define in an address pool, rather than an individual host route for each address. By default, the MAX adds dynamically assigned IP addresses to the routing table as individual host routes. To reduce the size of routing table advertisements, you can summarize the entire pool. The MAX routes packets to a valid host address, and rejects packets with an invalid host address.

Because the MAX creates a host route for every address assigned from the pools, and because host routes override subnet routes, the MAX correctly routes packets whose destination matches an assigned IP address from the pool. However, because the MAX advertises the entire pool as a route, and only knows privately which IP addresses in the pool are active, a remote network might improperly send the MAX a packet with an inactive IP address.

When the MAX receives a packet whose IP address matches an unused IP address in a pool, it either returns the packet to the sender with an ICMP reject message, or simply discards the packet. To enable the router to handle packets with destinations to invalid hosts on the summarized network, you must specify one of the following internal interfaces as the router:

Interface	Description
The reject interface (rj0)	The reject interface has an IP address of 127.0.0.2. When you specify this address as the router to the destination pool network, the MAX rejects packets to an invalid host on that network, appending an ICMP Host Unreachable message.
The black-hole interface (bh0)	The black-hole interface has an IP address of 127.0.0.3. When you specify this address as the router to the destination pool network, the MAX silently discards packets to an invalid host on that network.

See also *network alignment*, *route*, *router*.

POP—Point of Presence. A POP is the location of an Internet Service Provider's (ISP's) equipment. See also *ISP*.

port—A TCP/IP interface that defines the logical location in a computer where an application or process is running. When you define such a location, packets can reach an application from a remote system. There are certain well-known ports, such as port 21 used by FTP. Packet filters and firewalls make use of port addresses to restrict incoming and outgoing data and to secure an environment. The User Datagram Protocol (UDP) was developed to add the port address of an application or process to an IP packet, facilitating communication between applications over a network. See also *packet filter*, *firewall*, *IP*, *TCP/IP*, *UDP*.

POST—Power-On Self Test. A POST is a diagnostic test the MAX performs when it first starts up or after it completes a system reset. During a POST, the MAX checks system memory, configuration, installed cards, compression hardware, and T1 connections.

POTS—Plain Old Telephone Service. POTS denotes conventional analog voice transmission over telephone lines.

power interface—A MAX interface that accepts AC or DC power (depending on the model).

Power-On Self Test—See *POST*.

PPP—Point-to-Point Protocol. PPP provides a standard means of encapsulating data packets sent over a single-channel WAN link. It is the standard WAN encapsulation protocol for the interoperability of bridges and routers. PPP also allows direct dial-up access from a personal computer to a corporate LAN or Internet Service Provider (ISP). Using PPP ensures basic compatibility with non-Ascend devices. Both the dialing side and the answering side of the link must support PPP.

Figure 37 illustrates a single-channel PPP call.

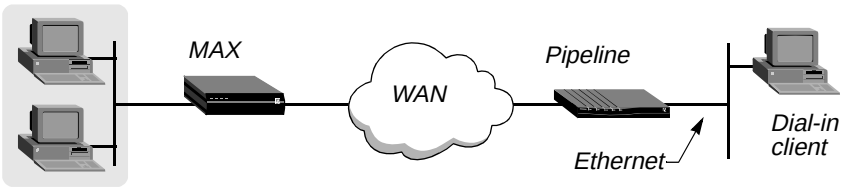

Figure 37. PPP connection

Typically, a dial-in device such as a modem or V.120 Terminal Adapter (TA) initiates a PPP session. The MAX unit's terminal-server software handles the call. If the terminal server detects a PPP packet from the caller, it passes the call on to the router, which handles it as a regular PPP connection. The caller never sees the terminal-server interface.

However, if the user's dial-in software does not support PPP, the user can still initiate a PPP session from within the terminal-server interface. To do so, a user can log into the terminal server in terminal mode and use the PPP command. Or, you can include the PPP command in an expect-send script.

During establishment of a PPP data link, the dialing and answering units exchange Link Control Protocol (LCP) packets to establish communications and configure the link. When the link is established, PPP provides for an optional authentication step before exchanging Network Control Protocols (NCPs).

See also *asynchronous PPP*, *expect-send script*, *ISP*, *LCP*, *NCP*, *router*, *synchronous PPP*, *terminal mode*, *terminal server*, *V.120 TA*.

PPTP—Point-to-Point Tunneling Protocol. PPTP is a protocol that enables you to extend your corporate network by means of private tunnels over the Internet. PPTP is a standard sponsored by Microsoft.

PPTP Network Server—See *PNS*.

precedence—In a Type-of-Service (TOS) policy or filter specification, the priority level of the data stream. The three most significant bits of the TOS byte are priority bits used to set the precedence for priority queuing. See also *TOS*, *TOS filter*.

preference—A way for the MAX to decide which route takes highest priority.

Routing Information Protocol (RIP) is a distance-vector protocol, which uses a hop count to select the shortest route to a destination network. Open Shortest Path First (OSPF) is a link-state protocol, which can take into account a variety of link conditions, such as the reliability or speed of the link, when determining the best path to a destination network. Because the metrics used by the two protocols are incompatible, the MAX supports route preferences.

By default, static routes and RIP routes have the same preference, so they compete equally. Internet Control Message Protocol (ICMP) Redirects take precedence over both, and OSPF takes precedence over everything. If a dynamic route's preference is lower than that of the static route, the dynamic route can temporarily hide a static route to the same network. However, dynamic routes age, and if no updates are received, they eventually expire. In that case, the hidden static route reappears in the routing table.

See also *dynamic route*, *hop count*, *ICMP*, *metric*, *OSPF*, *RIP*, *route*, *static route*.

Presentation layer—The second highest layer in the OSI Reference Model. The Presentation layer is responsible for presenting information in a format understandable to users and their applications. Data conversion, special graphics, compression, and encryption are some of the functions implemented at the Presentation layer. See also *OSI Reference Model*.

primary add-on number—For an ISDN BRI line, a number that enables the calling MAX to build multichannel calls. A multichannel call begins as a single-channel connection to one phone number. The calling unit then requests additional phone numbers it can dial to connect additional channels, and stores the add-on numbers it receives from the answering unit. When the MAX receives a multichannel AIM, BONDING, MP, or MP+ call, it reports the primary add-on number and the secondary add-on number to the calling party. The calling unit must integrate the add-on numbers with the phone number it dialed initially to add channels to the call. If you do not specify an add-on number and the calling MAX needs to add more channels, it redials the phone number it used to make the first connection. See also *AIM*, *BONDING*, *MP*, *MP+*, *secondary add-on number*.

primary DNS server—The first server the MAX attempts to access in order to perform name-address resolution on an IP network. If the primary DNS server is unavailable, the MAX attempts to use the secondary DNS server. See also *DNS*, *IP network*, *secondary DNS server*.

Primary Rate Interface line—See *E1 PRI line*, *T1 PRI line*.

primary WINS server—The first server the MAX attempts to access for WINS name-address resolution on a Telnet or raw TCP connection running under the MAX unit's terminal server interface. See also *secondary WINS server*, *WINS*.

PRI-to-T1 conversion—A feature that enables a single T1 PRI line to carry both data and voice traffic. A T1 PRI line offers reduced call-setup times, unrestricted 64-Kbps channels, call-by-call service selection, and enhanced data services such as Switched-384, Switched-1536, and MultiRate. With PRI-to-T1 conversion, any standard T1-based PBX can access voice circuits on the T1 PRI line, and LAN traffic can access both nailed-up and switched data circuits on the T1 PRI line. See also *D channel*, *MultiRate*, *nailed-up circuit*, *PCM*, *Switched-384*, *Switched-1536*, *switched circuit*, *T1 line*, *T1 PRI line*.

private circuit—See *nailed-up circuit.*

private-key encryption—An encryption method that uses a single key (that only the sender and receiver know) and a public encryption algorithm. Compare with *public-key encryption.* See also *encryption*.

private network—A network particular to an organization, and not connected to a public data network such as the Internet. See also *VPN*.

profile—A collection of settings that enable you to configure various aspects of an Ascend product. For example, a Connection profile enables you to specify the name, password, and network resources for a dial-in caller. See also *Connection profile*, *pseudo-user profile*, *subprofile*, *user profile*.

Programmable Read-Only Memory—See *PROM*.

PROM—Programmable Read-Only Memory. PROM is a memory chip on which the system can write data only once. A PROM chip retains its contents across power cycles and system resets. See also *EEPROM*.

promiscuous mode—A bridging mode in which the MAX unit's Ethernet controller accepts all packets and passes them up the protocol stack for a higher-level decision on whether to route, bridge, or reject them. Promiscuous mode is appropriate if you are using the MAX as a bridge. See also *bridge*.

protective ground—On a DB25 pin connector, the chassis ground connection between Data Terminal Equipment (DTE) and Data Circuit-Terminating Equipment (DCE). See also *DB25 pin connector*, *DCE*, *DTE*.

protocol—A set of rules governing message exchange over a network or internet. Examples of commonly used protocols are Transmission Control Protocol/Internet Protocol (TCP/IP), Point-to-Point Protocol (PPP), and Internetwork Packet Exchange (IPX). See also *internet*, *IPX*, *network*, *PPP*, *TCP/IP*.

Protocol Data Unit—See *PDU.*

Protocol Identifier—See *PID.*

proxy ARP—Proxy Address Resolution Protocol. Proxy ARP denotes a configuration in which one unit handles address resolution requests for another device. In an ARP request, a device asks a host to provide the host's physical address so that a connection can take place. ARP requests are broadcast only on the local network. If the MAX is the default router on a network and is configured in proxy mode, packets destined for any of the hosts on the network go to the MAX. If a remote host must respond to an ARP request, the MAX can respond on its behalf. See also *ARP*, *proxy mode*, *router*.

proxy mode—A mode in which a Connection profile assigns a local IP address to a remote host. Local hosts see the remote host as though it were on the local network. When calls are made to the remote host, the MAX acts on its behalf, replying to requests and forwarding packets. See also *proxy ARP*.

PSDN—Packet-Switched Data Network. A PSDN is a network in which no connection is required end-to-end. This type of network is very efficient for data transfer, and provides necessary redundancy. Other circuits are automatically available if a line goes down. See also *packet switching.*

pseudo-user profile—A RADIUS users file entry containing information that the MAX can query. Unlike a RADIUS user profile, it does not exist for the purpose of authenticating a user. Rather, it enables you to specify static route configurations, Frame Relay profile information, bridging entries, and other types of settings. See also *user profile*, *users file*.

PSPDN—Packet-Switched Public Data Network. A PSPDN is an X.25 network. See also *X.25*.

PSTN—Public Switched Telephone Network. A PTSN is a public circuit-switched network for telephone users. See also *circuit switching.*

public-key encryption—An encryption method that bases an encryption algorithm on the two halves of a long bit string. Each half of the bit sequence corresponds to a key. One key resides in a public-key library. Only a single party knows the other key. You can use either key to encrypt the data, but both keys are required to decrypt it. The sender can encrypt the data with the receiver's public key, and the receiver can decrypt it with the private key. Or, the sender can use private key to encrypt the message, and the receiver can use the public key to decrypt it. Compare with *private-key encryption.* See also *encryption*.

Public Switched Telephone Network—See *PSTN*.

Pulse Coded Modulation—See *PCM*.

PVC—Permanent Virtual Circuit. A PVC is a path maintained by two stations. The circuit is through the packet-switched network, but stays up all the time, regardless of whether or not data is on the line. Because the circuit is always up, there is no circuit setup time. Compare with *SVC*. See also *packet switching*.

Q

Q.931—An ITU (formerly CCITT) recommendation detailing the Layer Two protocol for an ISDN D-channel.

Q.931 en-bloc dialing—The process of sending all dialed digits to the MAX in one block. Q.931 en-bloc dialing is also known as *senderized digit transmission*.

Q.931 Layer 3 SETUP_ACK timer.—See *T302 timer*.

Q.931W GloBanD—See *GloBanD*.

Quality of Service—See *QoS*.

queue—A set of items arranged in a defined sequence. See also *backoff queue*, *RIP queue*, *SNMP queue*, *UDP queue*.

queue depth—The maximum number of unprocessed requests the MAX saves.

quiesce—To gracefully take a line or modem out of service.

QoS—Quality of Service. QoS encompasses the theory that one can measure, improve, and predict data rates, error rates, and other facets of network transmission. QoS is particularly important with relation to the transmission of high-bandwidth video and multimedia data. When you use the Resource Reservation Protocol (RSVP), you can expedite packets going through a gateway on the basis of criteria prepared in advance. Asynchronous Transfer Mode (ATM) also enables you to specify the level of quality you need. See also *ATM*, *RSVP*.

R

R2 signaling—An ITU-T standardized signaling protocol for establishing and clearing 64-kbps switched circuits on E1 digital trunks. Signaling is performed through a combination of A/B bit manipulation in channel 16 of the E1 frame, and inband MF tone generation and detection. The relevant specifications are found in ITU-T recommendations Q.400 to Q.490. R2 signaling is widely implemented in international markets in which ISDN PRI signaling is not yet available. See also *E1 line*, *ITU-T*.

RADIPAD—RADIUS IP Address Daemon. RADIPAD is a program that works with RADIUS to manage IP address pools centrally, so that connections can all acquire an address from a global pool, regardless of which system answers the call. RADIPAD runs on one RADIUS server on the network. A MAX sends an authentication request to RADIUS, and if the user profile contains an attribute to allocate an IP address from the global pool, RADIUS sends a request to RADIPAD to acquire the address. The MAX does not talk directly to RADIPAD, so it does not require additional configuration to use RADIPAD. See also *global IP address pool*, *IP address pool*, *RADIUS*, *RADIUS server*.

RADIUS—Remote Authentication Dial-In User Service. Using RADIUS, end users can have access to secure networks through a centrally managed server. RADIUS provides authentication for a variety of services, such as login, callback, Serial Line Internet Protocol (SLIP), and Point-to-Point Protocol (PPP). It also enables you to set up accounting. You can keep records on the number of packets the MAX transmitted and received, the protocol in use, the user name and IP address of the client, and other system information.

In Figure 38, the RADIUS server performs both authentication and accounting.

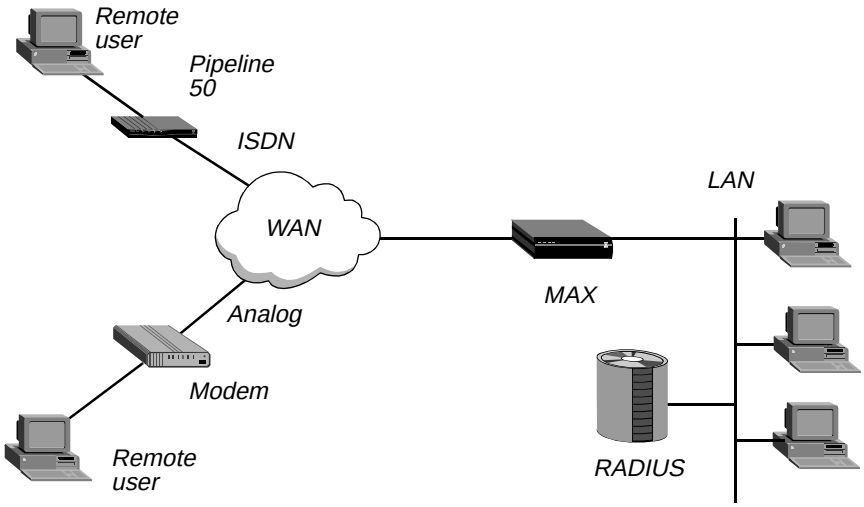

Figure 38. RADIUS authentication and accounting

In a RADIUS query, the MAX provides a user ID and password to the server. The server sends back a complete profile, which specifies routing, packet filtering, destination-specific static routes, and restrictions specific to the user. In addition, the MAX can use the data in the RADIUS database to create and advertise static routes, and to place outgoing calls.

The communications channel between a RADIUS client and server is provided by UDP/IP, with messages acknowledged. The primary advantage in using RADIUS to authenticate incoming calls is that you can maintain all user information offline on a separate UNIX-based server. The server can accept authentication requests from many machines, which makes swapping out one dial-in network server for another much easier.

See also *accounting*, *authentication*, *packet filter*, *pseudo-user profile*, *routing*, *static route*, *user profile*, *users file*.

RADIUS accounting—See *accounting*.

radiusd—The RADIUS daemon that runs with a flat ASCII `users` file.

RADIUS daemon—The daemon that provides users with RADIUS authentication and accounting. Ascend provides a RADIUS daemon that runs with a flat ASCII `users` file, and one that runs with a UNIX DBM database. The `radiusd` daemon runs with a flat file. The `radiusd.dbm` daemon runs with a UNIX DBM database.

radiusd.dbm—The RADIUS daemon that runs with a UNIX DBM database.

RADIUS IP Address Daemon—See *RADIPAD*.

RADIUS server—The machine on which the RADIUS daemon is running. A single RADIUS server can administer multiple security systems, maintaining profiles for thousands of users. See also *RADIUS, radiusd, RADIUS daemon, radiusd.dbm.*

RADIUS timeout—The number of seconds between retries to the RADIUS server. If the MAX is acting as a RADIUS client, it waits the specified number of seconds for a response to an authentication request. If it does not receive a response within that time, it times out and sends the authentication request to the next authentication server. See also *authentication server*, *RADIUS*.

RAI—Remote Alarm Indicator. An RAI indicates that a device on the T1 line is detecting framing-error conditions on the signal it receives. An RAI is also called a *Yellow Alarm signal*. See also *T1 line*.

RAM—Random Access Memory. RAM is computer memory that holds data temporarily. See also *DRAM*, *NVRAM*.

Random Access Memory—See *RAM*.

RARP—Reverse Address Resolution Protocol. RARP is a TCP/IP protocol that learns a workstation's hardware address and maps it to an IP address. See also *ARP*.

rate adaption—A data-transmission method that enables the MAX to send and receive data moving at a rate of 56 Kbps over a 64-Kbps channel. For incoming calls, the MAX automatically adapts the data received at 56 Kbps to the 64-Kbps channel. For outgoing calls, the MAX sets the data rate to either 64 Kbps or 56 Kbps.

For example, suppose a network consists of five switches, one of which uses a 56-Kbps line. The MAX sends data at 56 Kbps over the 64-Kbps line that connects the switch to the network. In doing so, the router drops one of the 8 bits of data and sends only 7 bits at a time.

V.110 and V.120 are both rate-adaption standards. See also *V.110*, *V.120*.

rate limit—See *multicast rate limit*.

Raw 802.3—See *802.3*.

Raw TCP—Raw Transmission Control Protocol. Raw TCP is a method of supporting encapsulation performed by an application that runs on top of TCP. Raw TCP must be understood by both the login host and the caller. As soon as the connection is authenticated, the MAX establishes a TCP connection to the host specified in the Connection profile or RADIUS user profile.

Raw TCP is also known as *TCP-Clear*. See also *TCP*, *V.120*.

Raw Transmission Control Protocol—See *Raw TCP*.

RBOC—Regional Bell Operating Company. An RBOC is one of seven companies created after the breakup of AT&T. The RBOCs are Ameritech, Bell Atlantic, Bell South, NYNEX, Pacific Telesis, Southwestern Bell, and U.S. West.

RD—Receive Data. RD is a signal that indicates that the modem is receiving data from a remote device. See also *DB25 pin connector*.

RDP—Reliable Data Protocol. RDP provides a reliable data transport service for packet-based applications. It is simple to implement, and works efficiently in environments that have long transmission delays and non-sequential delivery of message segments.

Read-Only Memory—See *ROM*.

real channels—Channels that connect directly to the MAX that owns the bundle in a stacked configuration. (Stacked channels connect to a MAX that transfers the data to or from the MAX that owns the bundle.)

For example, assume the initial call of an MP/MP+ bundle connects to MAX #1. This connection is a *real* channel. Next, the second call of the bundle connects to MAX #2. This connection is a *stacked* channel. MAX #1 is the bundle owner, and it manages the traffic for both channels of the bundle. MAX #2 forwards any traffic from the WAN to MAX #1, for distribution to the destination (Figure 39).

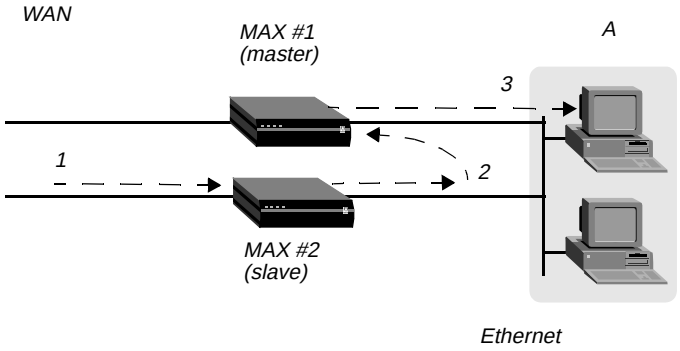

Figure 39. Packet flow from the slave channel to the Ethernet

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

Likewise, MAX #1 receives all Ethernet traffic destined for the bundle, and disperses the packets between itself and MAX #2 (Figure 40).

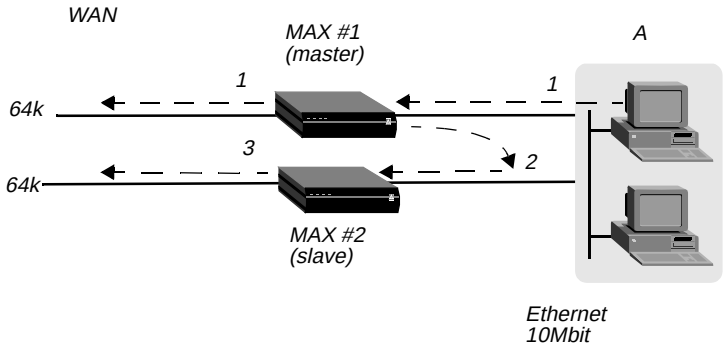

Figure 40. Packet flow from the Ethernet

MAX #1 forwards some of the packets across the WAN through a real channel. MAX #2 sends the rest of them through a stacked channel.

Compare with *stacked channels*. See also *bundle*, *bundle owner*, *stacks*.

Real-Time Transport Control Protocol—See *RTCP*.

Real-Time Transport Protocol—See *RTP*.

Receive Data—See *RD*.

receive data rate—The baud rate of data received by the MAX. Compare with *transmit data rate*.

receiver/transmitter echo—The acoustic echo generated at the calling endpoint of a connection. See also *echo cancellation, echo tail*.

Recognized Private Operating Agency—See *RPOA*.

Red Alarm signal—A signal indicating that an out-of-frame condition has lasted for more than 2.23 msec on the T1 line. See also *out-of-frame condition*, *T1 line*.

Reduced-Instruction-Set Computer—See *RISC*.

redundancy—A method of safeguarding against line and equipment failure during a transmission. Each method for transmitting signals has inherent error rates, and all physical media is subject to damage. In the event of hardware failure, a redundant line or unit can take over at any time. You should always have a redundant (backup) module for multiplexers and other critical equipment.

reject interface—An interface that enables the router to handle packets whose IP address matches an unused IP address in a summarized address pool. The reject interface has an IP address of 127.0.0.2. When you specify this address as the router to the destination pool network, the MAX rejects packets to an invalid host on that network, appending an ICMP Host Unreachable message. See also *pool summary*.

Reliable Data Protocol—See *RDP*.

Remote Alarm Indicator—See *RAI*.

Remote Authentication Dial-In User Service—See *RADIUS*.

remote device—A unit that resides across the WAN.

remote LAN Access—The process of allowing branch offices, telecommuters, and traveling computer users to access the corporate LAN backbone over digital or analog lines. The lines can be switched or nailed up. See also *analog line*, *digital line*, *nailed-up line*, *switched line*.

remote loopback—A procedure that tests the entire connection from host interface to host interface, enabling the MAX to place a call to itself over the WAN and to send a user-specified number of packets over the connection. The data loops at the AIM port interface of the remote MAX, and comes back to the local MAX. The remote loopback tests the MAX unit's ability to initiate and receive calls, and diagnoses whether the connection over the digital access line and the WAN is sound. Compare with *local loopback*. See also *analog loopback*, *digital loopback*, *loopback*.

remote management—A MAX management feature that uses bandwidth between sites over the management subchannel established by the Ascend Inverse Multiplexing (AIM) protocol. Any Ascend unit can control, configure, and obtain statistical and diagnostic information about any other Ascend unit. Multilevel security assures that unauthorized personnel do not have access to remote-management functions. See also *AIM*.

remote network—A network to which the MAX connects over the WAN.

Remote Port Module—See *RPM*.

Remote Procedure Call—See *RPC*.

remote profile—A user profile configured in RADIUS, TACACS, or TACACS+, as opposed to a Connection profile configured on the Ascend unit. See also *user profile*.

remote user—A user at a device not connected directly to the Ascend unit and not residing on the local Ethernet.

repeater—A device that receives, amplifies, and retransmits a signal, overcoming any attenuation in existence on the physical medium.

replay attack—A strategy for gaining illegal access to a system. During a replay attack, an unauthorized user records valid authentication information exchanged between systems, and then replays it later to gain entry. Token-card authentication protects your system against replay attacks. Because the token is a one-time-only password, replay is impossible. See also *token-card authentication*.

Request For Comments—See *RFC*.

Request To Send—See *RTS*.

Reset-Confirmation packet—On an X.25 connection, a packet sent in response to a Reset-Request packet. Once the sending device receives a Reset-Confirmation packet, it can send data on the logical channel. See also *logical channel, Reset-Request packet, X.25*.

Reset-Request packet—At the packet layer of an X.25 network, a packet that resets the packet-sequence number for the logical channel to 0 (zero), and removes any outstanding data and Interrupt packets from the Virtual Circuit (VC). See also *VC, X.25*.

Reset-Request timer—The number of ten-second ticks the MAX waits before retransmitting a Reset-Request packet on an X.25 network. See also *Reset-Request packet, X.25*.

Reset Retries—The number of times the MAX retransmits a Reset-Request packet on an X.25 network before clearing a call. See also *Reset-Request packet, X.25*.

Resource Reservation Protocol—See *RSVP*.

Response timer—On an X.25/T3POS connection, a timer that indicates the amount of time the Packet Assembler/Disassembler (PAD) waits for a SYN signal from the Data Terminal Equipment (DTE). The SYN signal indicates that the response from the DTE is being delayed and that the link is still alive. See also *DTE, PAD*.

Restart-Confirmation packet—On an X.25 network, a packet that signals a sending device that it can again issue a call to establish a Virtual Circuit (VC). See also *VC, X.25*.

Restart-Request packet—At the packet layer of an X.25 network, a packet that clears all Virtual Circuits (VCs). See also *Restart Retries, Restart timer, X.25*.

Restart Retries—The number of times the MAX transmits a Restart-Request packet before waiting indefinitely for a response. See also *Restart-Request packet, X.25*.

Restart timer—On an X.25 network, the number of ten-second ticks the MAX waits before retransmitting a Restart-Request packet. See also *Restart-Request packet, X.25*.

restricted load—A software load containing only essential system software and not meant to be run in a working environment. A restricted load does not have full functionality and is to be used only to upgrade to an extended load. Restricted loads *do* allow you to access the unit by means of Telnet. Compare with *extended load, fat load, thin load*.

retry limit—The maximum number of times the MAX sends a packet or frame, or attempts to connect to another device, before giving up and clearing the connection. See also *retry timeout*.

retry timeout— The number of seconds between retries. See also *retry limit*.

Reverse Address Resolution Protocol—See *RARP*.

RFC—Request for Comments. RFC denotes the document series, begun in 1969, that describes the Internet suite of protocols and related experiments. Not all RFCs describe Internet standards, but all Internet standards are written up as RFCs. The RFC series of documents is unusual in that the proposed protocols are distributed by the Internet research and development community, acting on its own behalf. The protocols do not go through the formal review and standardization process promoted by organizations such as ANSI. A complete list of RFCs resides at http://www.internic.net/rfc/.

RI—Ring Indicate. RI is a signal that indicates that a call is coming into a unit.

ringback tone—A tone that the MAX generates when it answers an analog modem call. The calling modem hears the ringback tone, and then begins the modem protocol. See also *modem*.

Ring Indicate—See *RI*.

RIP—Routing Information Protocol. RIP is a distance-vector protocol found in both the NetWare and TCP/IP protocol suites. The protocol keeps a database of routing information that it gathers from periodic broadcasts by each router on a network.

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

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

The MAX recognizes network number –2 (0xFFFFFFFE) as the IPX RIP default route. When it receives a packet for an unknown destination, the MAX forwards the packet to the IPX router advertising the default route. If more than one IPX router is advertising the default route, the unit makes a routing decision based on the hop and tick count. For example, if the MAX receives an IPX packet destined for network 77777777 and it does not have a RIP table entry for that destination, the MAX forwards the packet towards network number FFFFFFFE, if available, instead of simply dropping the packet.

See also *default route*, *distance-vector metric*, *hop*, *IPX*, *IPX router*, *router*, *routing*, *split horizon*, *TCP/IP*, *tick*.

RIP queue—A queue containing unprocessed Routing Information Protocol (RIP) requests. Compare with *backoff queue*, *SNMP queue*, *UDP queue*. See also *queue*.

RISC—Reduced-Instruction-Set Computer. RISC is a microprocessor for decoding and executing a small set of instructions, thereby optimizing performance.

Rlogin—An Application-layer, remote-login service provided by Berkeley UNIX. On the MAX, Rlogin is available only from an asynchronous dialup session to the terminal server. See also *Application layer*.

robbed-bit signaling—See *inband signaling*.

ROM—Read-Only Memory. ROM is computer memory whose contents can be read and executed, but not modified. See also *EEPROM*, *PROM*.

rotary hunt group—A type of hunt group in which the incoming call hunts on a rotating basis for an available channel to ring and answer the call. See also *hunt group*.

route—The path that data takes from its source network to its destination network. See also *IP route, IPX route*.

router—A device that determines a path from a host on one network to a host on another. The networks may be in close proximity, or may be separated by long distances. A router has access to the three lowest OSI layers, and generally operates at the Network layer. To route a packet, a router uses the logical address specified as the packet's destination field, and determines the next router (if any) to which the packet must travel to reach its destination. All routers share information about the current topology and state of the network, maintaining routing tables that reflect the latest information. See also *IP router*.

Router Home Agent—In an Ascend Tunnel Management Protocol (ATMP) configuration, a Home Agent whose routing module forwards packets it receives from the Foreign Agent onto the local network. The network can be the Home Network, or it can support another router that can connect to the Home Network. In either case, packet delivery relies on a routing mechanism, such as a static or dynamic route, and not on a WAN connection. Compare with *Gateway Home Agent*. See also *ATMP, dynamic route, Foreign Agent, Home Agent, Home Network, static IP route, static IPX route*.

routing—A method of determining how to forward a data packet to the proper destination. See also *IP routing, IPX routing, OSPF, RIP, route, router*.

Routing Information Protocol—See *RIP*.

routing table—See *IP routing table*.

Routing Table Maintenance Protocol—See *RTMP*.

RPC—Remote Procedure Call. An RPC is a method in which a program on one device can transparently use a procedure on another device. RPCs are often used in client-server architectures.

RPM—Remote Port Module. An RPM is an Ascend unit that enables you to extend data, signaling, and control ports to local applications over Unshielded Twisted Pair (UTP) cable. See also *UTP cable*.

RPOA—Recognized Private Operating Agency. An RPOA is an operating agency that runs a telecommunications service.

RS-232—An EIA standard that specifies various electrical and mechanical characteristics for interfaces between Data Terminal Equipment (DTE) and Data Circuit-terminating Equipment (DCE) devices. The standard applies to both synchronous and asynchronous binary data transmission at rates below 64 Kbps. RS-232 is also known as *EIA/TIA-232*. Compare with *RS-422, RS-423*. See also *asynchronous transmission, DCE, DTE, synchronous transmission*.

RS-232C—The latest version of the RS-232 standard. See also *RS-232*.

RS-366—An EIA standard that specifies dialing commands to network-access equipment. Although RS-366 uses RS-232 electrical specifications, it specifies different pinouts and signal functions. See also *RS-232*.

RS-422—A standard EIA interface for connecting serial devices. Along with RS-423, RS-422 replaces the RS-232 standard. RS-422 supports higher data rates than RS-232, and offers greater protection against electrical interference. RS-422 supports multipoint connections, while RS-423 does not. RS-422 and RS-423 are often referred to collectively as *EIA-530*. Compare with *RS-232*, *RS-423*. See also *EIA*, *multipoint link.*

RS-423—A standard EIA interface for connecting serial devices. Along with RS-422, RS-423 replaces the RS-232 standard. RS-423 supports higher data rates than RS-232, and offers greater protection against electrical interference. Unlike RS-422, RS-423 supports only point-to-point connections. RS-422 and RS-423 are often referred to collectively as *EIA-530*. Compare with *RS-232*, *RS-422*. See also *EIA*, *point-to-point link*.

RS-449—A standard EIA Physical-layer interface. RS-449 is a a faster version of RS-232, and allows longer cable extension. RS-449 can run at speeds of up to 2 Mbps, and is also known as *EIA/TIA-449*. Compare with *RS-232*.

RSVP—Resource Reservation Protocol. RSVP enables a router to reserve bandwidth for time-sensitive data transmissions, resulting in smooth reception of voice and video data. A client can use RSVP to request that a router set aside a certain amount of bandwidth to handle the incoming call. If sufficient bandwidth does not exist, the request enters a queue and remains there until the appropriate amount of bandwidth becomes available.

RTCP—Real-Time Transport Control Protocol. RTCP enables a system to monitor the Quality-of-Service (Qos) and to transmit information about the participants in a real-time audio or video session. See also *RTP*.

RTMP—Routing Table Maintenance Protocol. Apple Computer's proprietary routing protocol.

RTP—Real-Time Transport Protocol. RTP provides end-to-end delivery services for real-time data, such as interactive video and audio. RTP identifies the type of data being transmitted and provides packet sequencing, timestamping, and monitoring services. Using RTP, a unit can transmit data using multicast services if the network provides them. RTP does not provide Quality-of-Service (QoS) guarantees, nor does it ensure reliable delivery with all packets in the proper sequence. See also *RTCP*.

RTS—Request To Send. RTS is a signal that a transmitter sends to a receiver in order to indicate that it wants to begin sending data. If the receiver is ready for the transmission, it responds with a Clear To Send (CTS) signal. See also *CTS.*

S

SafeWord AS—See *Enigma Logic SafeWord server*.

SafeWord authentication—A form of token-card authentication in which RADIUS forwards a connection request to an Enigma Logic Safeword server, a type of authentication server. The SafeWord server sends an Access-Challenge packet back through the RADIUS server and the MAX to the user dialing in. The user sees the challenge message, obtains the current password from his or her token card, and enters the current password (called a *token*).

The token travels back through the MAX and the RADIUS server to the SafeWord server. The SafeWord server sends a response to the RADIUS server, specifying whether the user has entered the proper user name and token. If the user enters an incorrect token, the SafeWord server returns another challenge, and the user can again attempt to enter the correct token. The server sends up to three challenges. After three incorrect entries, the MAX terminates the call.

See also *authentication, authentication server, RADIUS server, SafeWord token, token-card authentication, token-card server*.

SafeWord token—A randomly generated access code that a user obtains from a token card when using a SafeWord authentication server. See also *SafeWord authentication.*

SAM—Secure Access Manager. SAM gives you a high degree of centralized control over the security functions of an entire network. Through this Windows-based application, you can configure Secure Access Firewall(s) offline, and download the configuration to remote locations. See also *Secure Access Firewall*.

SAP—Service Access Point. A SAP is a defined location through which a procedure at one OSI layer can provide services to the next layer above it. Each SAP has a unique address in hexadecimal format. See also *DSAP, OSI Reference Model, SSAP*.

SAP—Service Advertising Protocol. SAP is a NetWare protocol that operates at the Transport layer and enables servers to inform other devices about the services they have available. Each server advertises its services using a SAP packet. Each router on the network retrieves the SAP packets and builds a database of all the servers it knows about. The router then broadcasts this information to other routers, either at a set interval or whenever the database changes.

The MAX follows standard SAP behavior for routers when connecting to non-Ascend units across the WAN. However, when it connects to another Ascend unit configured for IPX routing, both ends of the connection exchange their entire SAP tables, so all remote services are immediately added to the MAX unit's SAP table.

When a NetWare client sends a SAP request to locate a service, the MAX consults its SAP table and replies with its own hardware address and the internal network address of the requested server. This behavior is analogous to the use of proxy ARP in an IP environment. The client can then transmit packets whose destination address is the internal address of the server. When the MAX receives those packets, it consults its Routing Information Protocol (RIP) table. If it finds an entry for that destination address, it brings up the connection (unless it is already up) and forwards the packet. See also *IPX router, IPX server, proxy ARP, RIP*.

SAP filter—Service Advertising Protocol filter. A filter that determines which SAP advertisements the MAX forwards or drops.

The router examines incoming and outgoing SAP packets to see whether certain fields in the packet match the filter. The MAX applies input filters to all SAP packets it receives. Input filters screen advertised services and exclude them from (or include them in) the MAX service table. The MAX applies output filters to SAP response packets it transmits. If it receives a SAP request packet, the MAX applies output filters before transmitting the SAP response, and excludes services from (or includes them in) the response packet.

A SAP filter enables you to control the size of resident SAP tables and reduce bandwidth usage. You can also use a SAP filter to restrict a user's view of services on the network. By turning off IPX SAP, you can prevent the MAX from sending its SAP table or from receiving a remote site's SAP table.

See also *SAP*.

SAP home server proxy—Service Advertising Protocol home server proxy. SAP home server proxy is a feature that enables you to give remote users access to the same resources as local users. Rather than relying on the built-in functionality of SAP, you can configure the MAX to direct SAP broadcasts to specified networks. The SAP responses come from servers on these networks, rather than from servers near the MAX.

SDRP—Source Demand Routing Protocol. SDRP supports source-initiated selection of interdomain routes, working along with the intermediate node selection provided by Border Gateway Protocol (BGP) and Inter-Domain Routing Protocol (IDRP). See also *BGP*, *IDRP*.

sealing—The ability of the IDSL card to send some current (40V) on the line. You typically use this feature to keep the physical connection from corroding. Corrosion can occur when no activity occurs on the line. See also *IDSL card*.

secondary add-on number—For an ISDN BRI line, an alternate number that enables the calling MAX to build multichannel calls. See also *AIM*, *BONDING*, *MP*, *MP+*, *primary add-on number*.

secondary DNS server—The second server the MAX attempts to access in order to perform name-address resolution on an IP network. See also *DNS*, *IP network*, *primary DNS server*.

secondary WINS server—The second server the MAX attempts to access for WINS name-address resolution on a Telnet or raw TCP connection running under the MAX unit's terminal server interface. See also *primary WINS server*, *WINS*.

Secure Access Firewall—An Ascend software option that stops intruders from breaking and entering into your network. A firewall is similar to a filter, but is more complex, dynamically changing in response to the characteristics of the packets that pass through it. The firewall affects which packets are allowed to reach the network, and which packets can leave the network for another interface. Typically, you can design a firewall to flag a packet with specific bit patterns, and put rules into action that cause other rules to be created. For a firewall to take effect, you must apply it to a LAN or WAN interface. See also *filter persistence*, *SAM*.

Secure Access Manager—See *SAM*.

SecureConnect—Ascend's family of security products. These products provide a complete solution for setting up and using secure sessions across the Internet. They provide companies with the industry's most robust solution available for safeguarding information. The SecureConnect family of products includes:

- Ascend Access Control
- Secure Access Firewall

See also *Ascend Access Control*, *Secure Access Firewall*.

SecurID—A proprietary brand of token card used with a Security Dynamics ACE/Server. The server generates a code based on a user's ID, a password, and specific information encoded in the card. When the user attempts to log into a secure network, the token-card server requests a code generated within the previous 60 seconds. The server interprets the code. If it is genuine, the server grants access to the user. See also *ACE authentication*, *ACE token*, *authentication*, *authentication server*, *Security Dynamics ACE/Server*, *token card*, *token-card authentication*, *token-card server*.

SecurID authentication—See *ACE authentication*.

SecurID server—See *Security Dynamics ACE/Server*.

security card—See *token card*.

security-card authentication—See *token-card authentication.*

security-card server—See *token-card server.*

Security Dynamics ACE/Server—A type of authentication server that performs token-card authentication. A Security Dynamics ACE/Server is also known as a *SecurID server.* See also *ACE authentication*, *ACE token*, *authentication*, *authentication server*, *SecurID*, *token card*, *token-card authentication*, *token-card server.*

Security profile—A profile that consists of parameters you can set to control access to the MAX.

seed router—An IPX or AppleTalk router from which other routers learn their network configurations. Compare with *nonseed router*. See also *AppleTalk routing*, *IPX router*.

senderized digit transmission—See *Q.931 en-bloc dialing*.

Sequenced Packet Exchange—See *SPX*.

serial communication—Communication through the serial port of a device. For Windows 3.1, the maximum speed of the serial port is 19,200. For Windows 95, the serial port limit is 921,600. These limitations are subject to change with the development of a faster serial bus. See also *serial port*, *serial transmission.*

serial connection—A link between the serial ports of two devices. See also *serial communication*, *serial port*, *serial transmission*.

serial host—A device (such as a videoconferencing codec) that is connected to a serial WAN port communicating over a point-to-point link. To a serial host, the MAX appears to be a cable or Data Circuit-terminating Equipment (DCE). See also *codec*, *DCE*, *point-to-point link*, *serial WAN port*.

Serial Line Internet Protocol—See *SLIP*.

serial port—A port that transmits and receives asynchronous or synchronous serial data. See also *asynchronous transmission*, *serial transmission*, *synchronous transmission*.

serial transmission—A form of data transmission in which only one line carries all eight bits of a byte. In serial transmission, one bit follows another (as opposed to parallel transmission, in which the bits travel simultaneously, each on a different wire). Serial transmission can be either synchronous or asynchronous. Synchronous communication requires additional lines for transmitting handshake or timing signals. In asynchronous communication, the data itself contains synchronization information, so neither handshake nor clock signals are necessary. See also *asynchronous transmission*, *synchronous transmission*.

serial V.35 DTE port—See *serial WAN port*.

serial WAN port—A port that provides a V.35/RS-449/X.21 WAN interface, typically used to connect the MAX to a Frame Relay switch.The clock speed received from the link determines the serial WAN data rate. The maximum acceptable clock is 8 Mbps. The clock speed at the serial WAN port has no effect on the bandwidth of other WAN interfaces in the MAX. The serial WAN port is also called a *serial V.35 DTE port*. See also *serial transmission*.

Series56 Digital Modem module—A MAX module that supports 8, 12, or 16 K56flex-compatible digital modems. Each digital modem provides analog and cellular connections at rates of up to 56 Kbps. When you install the card, a remote user can dial into the MAX over a T1 line. You can install a maximum of 72 digital modems in the MAX. The Series56 Digital Modem module supports Rockwell 2.084 firmware. Rockwell 2.084 firmware supports V.90, K56flex, K56plus, and all slower, standard modem speeds. See also *analog data*, *digital modem*, *E1 PRI line*, *K56flex*, *T1 line*, *V.90*.

server—A device or program that provides services to hosts on a network.

server key—A RADIUS key used to validate the authenticator field on requests and generate the authenticator on responses. See also *authenticator field*, *RADIUS*.

Service Access Point—See *SAP*.

Service Advertising Protocol—See *SAP*.

service management—The ability to create and control data services, enabling fast and cost-effective deployment and ongoing management. For service providers, service management means the ability to create a variety of services targeted for varied enterprise market segments with appropriate traffic management plans and price models. For enterprises, service management means new, public data services targeted and priced for business needs that the enterprise can measure directly with access to network information. See also *Navis*, *NavisAccess*, *NavisCore*, *NavisXtend*.

Service Profile Identifier—See *SPID*.

session—The state a connection reaches when two parties can communicate with each other.

session ID—A unique ID that denotes a particular MAX session. The MAX can pass a session ID to SNMP, RADIUS, or other external entities. See also *session*, *session ID base*.

session ID base—The base number for calculating a session ID. If the value of the session ID base is nonzero, the MAX uses it as the initial base for calculating session IDs after a system reset. The system increments the ID for each subsequent session by 1. If the session ID base is zero, the MAX sets the initial base for session IDs to the absolute clock. For example, if the clock is 0x11cf4959, the subsequent session IDs uses 0x11cf4959 as a base. However, if the clock changes and the system reboots or clears NVRAM, session IDs may be duplicated. See also *session*, *session ID*.

session key—In RADIUS, a key that associates the client request with the user session. See also *RADIUS*.

Session layer—The third highest layer in the OSI Reference Model. The Session layer synchronizes the data in a network connection, maintains the link until the transmission is complete, handles security, and makes sure that the data arrives in the proper sequence. Gateway communications are implemented at the Session layer. Examples of Session-layer protocols are AppleTalk Data Stream Protocol (ADSP), NetBEUI (an extension of NetBIOS), NetBIOS, and Printer Access Protocol (PAP). See also *OSI Reference Model*.

Session Status window—A status window that displays the number of active bridging/routing and modem (terminal-server) sessions on the Ethernet. Following is a sample Session Status window:

```
|--------------------|
|90-100 Sessions     |
|>  1 Active         |
| 0 slc-lab-236      |
|                    |
|--------------------|
```

When this window is active, you can scroll down to see the name, address, or CLID of each connected device. Each line starts with a one-character session-status indicator. For example, 0 means *online*. For terminal-server sessions, the window identifies the modem number. See also *status window*.

Setup message—See *ISDN Call Setup message*.

SG—Signal Ground. A connection to which the system refers all electrical signals on an interface. See also *DB25 pin connector*.

shared secret—A password shared between the MAX and the RADIUS server. See also *RADIUS server*.

Shielded Twisted Pair cable—See *STP cable*.

Signal Ground—See *SG*.

Signaling System 7—A protocol architecture that specifies a series of Signaling Points (SPs) and Signaling Transfer Points (STPs) connected on a network. The SPs are hosts from which signaling messages originate and terminate. The STPs are packet switches that perform message routing between adjacent SPs or STPS. The Network Services Part (NSP) of the Signaling System 7 provides reliable message transfer, and corresponds to the lower three layers of the OSI model. The NSP consists of a Message Transfer Part (MTP) and a Signalling Connection Control Part (SCCP). See also *OSI Reference Model*.

signaling types—A mutually agreed-upon way to maintain synchronization and transfer data effectively between endpoints. The sending device and the receiving device must send signals in order to synchronize their clocks and determine where one block of data ends and the next begins. Inband signaling, ISDN D-channel signaling, and Non-Facility Associated Signaling (NFAS) are all examples of signaling types. See also *inband signaling, ISDN D-channel signaling, NFAS*.

Simple Mail Transfer Protocol—See *SMTP*.

Simple Network Management Protocol—See *SNMP*.

Simple Network Time Protocol—See *SNTP*.

single-address NAT—Single-address Network Address Translation. Single-address NAT provides a method of translating an address for hosts on the local network. without borrowing IP addresses from a DHCP server on the remote network.

For incoming calls, the MAX can perform NAT for multiple hosts on the local network using its own IP address. The MAX routes incoming packets for up to 10 different TCP or UDP ports to specific servers on the local network. Translations between the local network and the Internet or remote network are static and need to be preconfigured.

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

Compare with *Multiple-address NAT*. See also *NAT for LAN*.

S interface—See *S/T interface*.

SLIP—Serial Line Internet Protocol. SLIP enables your computer to send and receive IP packets over a serial link. The MAX does not support a direct SLIP dial-in, because SLIP does not support authentication. However, if SLIP is enabled in the terminal server, a user can initiate a SLIP session, and then run an application such as File Transfer Protocol (FTP). To begin a SLIP session, the user can log into the terminal server in terminal mode and use the SLIP command. Or, you can include the SLIP command in an expect-send script. Compare with *CSLIP*. See also *expect-send script, FTP, terminal mode, terminal server*.

slot—On the backplane of a MAX, the connector that provides the physical and electrical connection between a card and the MAX unit's base resources.

slot card—A card you install on the MAX in order to enhance its functionality. For example, you can install a digital modem card to provide digital modem access.

slot compression—Compression in which the slot ID does not appear in any VJ-compressed packet but the first packet in the data stream. When you turn on VJ compression, the MAX removes the TCP/IP header, and associates a TCP/IP packet with a connection by giving it a slot ID. The first packet coming into a connection must have a slot ID, but succeeding packets need not have one. By default, the MAX uses slot compression: if the packet does not have a slot ID, the MAX associates it with the last-used slot ID. See also *VJ compression*.

SMDS—Switched Multimegabit Data Service. SMDS is a packet-based service that enables the creation of high-speed data networks with rates of up to 45 Mbps.

SMTP—Simple Mail Transfer Protocol. In the TCP/IP protocol suite, SMTP is an Application-layer protocol that uses the TCP Transport-layer protocol to send and receive electronic mail. See also *TCP/IP*.

SNAP—SubNetwork Access Protocol. SNAP is a protocol specification for the format of the Media Access Control (MAC) header of an IPX frame. SNAP includes the IEEE 802.3 protocol format plus additional information in the MAC header. Compare with *802.2*, *802.3*, *Ethernet II*. See also *IPX frame*, *MAC*.

SNMP—Simple Network Management Protocol. SNMP is a standard way for computers to share networking information.

In SNMP, two types of communicating devices exist: agents and managers. An agent provides networking information to a manager application running on another computer. The agent can be polled by the manager, and can also use a message called a traps-PDU to send unsolicited information to the manager when an unusual event occurs. The MAX is an example of an SNMP agent. The agents and managers share a database of information, called the Management Information Base (MIB).

The MAX supports SNMP MIB II, T1 MIB, and Ascend Enterprise MIBs. A manager that uses the Ascend Enterprise MIB can query the MAX, set parameters, sound alarms when certain conditions appear, and perform other administrative tasks. You can therefore manage the MAX from a central SNMP manager, such as SunNet Manage or HP Open View.

SNMP security uses the community name that the manager sends (with each polling request) and that the agent sends (with each traps-PDU). Ascend supports two community names: one with read-only access, and the other with read/write access to the MIB.

SNMP queue—A queue containing unprocessed Simple Network Management Protocol (SNMP) requests. Compare with *backoff queue*, *RIP queue*, *UDP queue*. See also *queue*.

SNTP—Simple Network Time Protocol. SNTP is a protocol that enables a group of servers to synchronize their clocks with reference to a primary time server. See also *SNTP server*, *UTC*.

SNTP server—A server that retrieves the correct time from an official source and distributes the information to other servers and networks. See also *SNTP*, *UTC*.

socket—A TCP/IP interface that facilitates a two-way link between systems, enabling applications to run over a connectionless network. A socket is defined by two addresses: the IP address of the host computer, and the port address of the application or process running on the host. See also *IP address*, *port*, *TCP/IP*.

socket number—A unique value assigned to a socket in a network. See also *socket*.

software compression—See *compression*.

software flow control—A method of flow control that uses the special characters XON and XOFF in the data stream. Compare with *hardware flow control*. See also *flow control*.

software handshaking—A synchronization method that uses the XON and XOFF characters to signal the beginning and end of a transmission. XON indicates that the device can receive data. XOFF interrupts the flow of data until an XON is sent. Compare with *hardware handshaking*. See also *handshaking*.

source address—In a frame, packet, or message sent over a bridged or routed connection, the IP, IPX, AppleTalk, or hardware address of the device that sent the transmission. Compare with *destination address*. See also *AppleTalk routing*, *hardware address*, *IP address*, *IP routing*, *IPX bridging*, *IPX routing*.

Source Demand Routing Protocol—See *SDRP*.

source port—The port from which a transmission originates, such as a User Datagram Protocol (UDP) port on an authentication server, or an Simple Mail Transfer Protocol (SMTP) port on a mail server. Compare with *destination port*. See also *SMTP*, *UDP*, *UDP port*.

Source Service Access Point—See *SSAP*.

space parity—A synonym for *even parity.* See also *parity*.

spanning—See *call spanning*.

SPID—Service Profile Identifier. A SPID is a number that the telephone company uses at the Central Office (CO) switch to identify services on your ISDN line. Each SPID is derived from a telephone number.

split horizon—An IPX mechanism for preventing circular routes and reducing network traffic. The simple split-horizon scheme omits routes learned from one neighbor in updates sent to that neighbor. A split horizon with poison reverse policy includes such routes in updates, but sets each metric to infinity.

spoofing—A procedure that enables a device to a) mimic a legitimate network host and gain access to data within a private network, causing a security breach, b) receive an IP address from a DHCP server across a slow WAN link., or c) imitate a return watchdog packet for the purpose of allowing clients to remain logged into a remote server. See also *DHCP spoofing*, *IP address spoofing*, *IPX spoofing*, *SPX spoofing*, *watchdog spoofing*.

SPX—Sequenced Packet Exchange. SPX is a Transport-level protocol that enables a system to perform connection-oriented packet delivery. See also *IPX*, *SPX spoofing*.

SPX spoofing—A feature that allows the WAN connection to remain idle while the application(s) requiring it are idle.

NetWare applications that require guaranteed packet delivery, such as Print Server (PSERVER), Remote Printer (RPRINTER), and Remote Console (RCONSOLE), use the NetWare SPX protocol. The client's SPX watchdog process monitors the connection with the server while the connection is idle. While an SPX application is running, the SPX watchdog sends a query that brings up the WAN connection every 14 seconds. To allow Netware SPX clients to remain logged in without keeping the WAN connection up in times of inactivity, the Ascend unit responds to SPX watchdog requests from the LAN with a fake SPX-watchdog-reply packet, and drops any SPX-watchdog-alive packets from the LAN without sending them on to the WAN.

Compare with *DHCP spoofing*, *IP address spoofing*, *IPX spoofing*, *watchdog spoofing*. See also *SPX*.

SSAP—Source Service Access Point. An SSAP is the Service Access Point (SAP) address at which at a Network-layer procedure requests services from the Logical Link Control (LLC) layer. See also *DSAP*, *SAP*.

Stac Lempel-Ziv standard compression—See *Stac LZS compression*.

Stac compression—On the MAX, a compression option that specifies an Ascend-modified version of draft 0 of the CCP (Compression Control Protocol). The Stac option is an Ascend variant of the Stac LZS compression method. It was implemented before Stac LZS was standardized. Compare with *Stac LZS compression.*

Stac-9 compression—On the MAX, a compression option that indicates the method specified by draft 9 of the Stac LZS compression protocol. Compare with *Stac compression*. See also *Stac LZS compression*.

Stac LZS compression—Stac Lempel-Ziv standard compression. Developed by Stac Incorporated, Stac LZS compression can triple data rates. Compare with *Stac compression*. See also *Stac-9 compression*.

stacked channels—Channels that connect to a MAX that transfers the data to or from the MAX that owns the bundle in a stacked configuration. For more information, see the description of *real channels.*

stacks—A group of MAX units that act as a single, logical unit with a single stack name. A stack allows incoming Multilink Protocol (MP) or Multilink Protocol Plus (MP+) calls to span multiple MAX units on a single LAN. There is no master unit in a stack. A MAX can become a member of a stack or leave a stack at any time, and there is no requirement to join a stack. MAX units in a stack find each other using an Ethernet multicast packet. Because multicast packets are unlikely to cross a router, all members of a stack must reside on the same physical LAN. Each MAX must use the same version of the software. See also *MP, MP+, multicast.*

start bit—In asynchronous transmission, a bit that indicates the beginning of a new character. It is always 0 (zero). Compare with *parity bit, stop bit.* See also *asynchronous transmission*.

Start record—A RADIUS-accounting or call-logging record that contains information about the beginning of a session with the MAX. See also *Start session.*

Start session—An event denoting the beginning of a session with the MAX. Information about a Start session event appears in a RADIUS-accounting or call-logging Start record.

static IP route—A path that specifies a destination IP network and the gateway (next-hop router) to get to that network. For example, if a Connection profile specifies the destination address of a host on a remote subnet, but the packets must be routed through an intermediary device to reach that host (and RIP or OSPF is not enabled), you must configure a static route specifying the intermediary device as the gateway. Figure 41 shows an example.

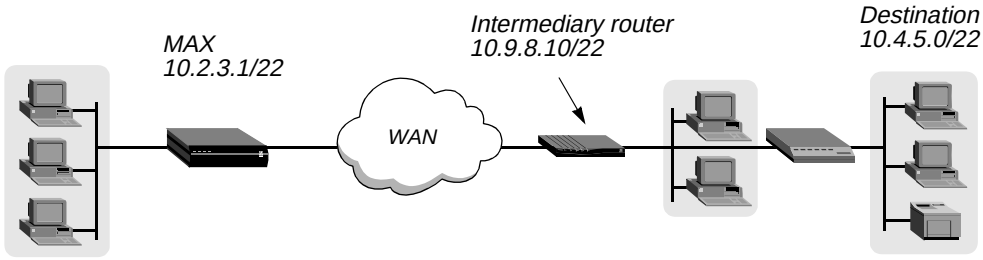

Figure 41. Static route to a remote subnet

Compare with *dynamic route*, *multipath route*. See also *Connection profile*, *IP address*, *IP network*, *pseudo-user profile*, *user profile*.

static IPX route—A route that contains all the information necessary to reach one IPX server on a remote network. The MAX adds the static routes upon initialization. When the MAX receives an outgoing packet for a server, it finds the corresponding profile and dials the connection. You must manually update static routes whenever the administrator at the remote end removes the specified server or updates its address. You do not need to create IPX routes to servers that reside on the local Ethernet network. See also *IPX server*, *pseudo-user profile.*

static password—A password specified in a Connection profile or RADIUS user profile. The user must enter the password to gain access to the MAX. See also *Connection profile*, *user profile*.

static route—See *static IP route*, *static IPX route*.

station—See *host*.

status window—Any one of eight windows displayed on the right side of the screen in the MAX configuration interface. These status windows provide a great deal of read-only information about what is currently happening in the MAX. Following are the status windows:

```
|--------------------| |--------------------|
|10-100 1234567890   | |10-200 1234567890   |
| L1/LA nnnnnnnnnn   | | L2/RA ..........   |
|   12345678901234   | |   12345678901234   |
|   nnnnnnnnnnnnnn   | |   ..............   |
|--------------------| l--------------------|
|90-100 Sessions     | |00-200 15:10:34     |
|>  1 Active         | |>M31  Line    Ch    |
| O slc-lab-236      | | LAN session up     |
|                    | | slc-lab-236        |
|--------------------| |--------------------|
|90-300 WAN Stat     | |90-400 Ether Stat   |
|>Rx Pkt:     184318^| |>Rx Pkt:    3486092 |
| Tx Pkt:     159232 | | Tx Pkt:      10056 |
|    CRC:          0v| |    Col:       3530 |
|--------------------| |--------------------|
|00-100 Sys Option   | |Main Status Menu    |
|>Security Prof: 1  ^| |>00-000 System     ^|
| Software +5.0A0+   | | 10-000 Net/T1      |
| S/N: 5210003      v| | 20-000 Net/T1     v|
|--------------------| |--------------------|
```

The displays contains the following windows:

- Line Status (the two topmost windows)
- Session Status (the second window on the left)
- System Status (the second window on the right)
- WAN Status window (the third window on the left)
- Ethernet Status (the third window on the right)
- Sys Option (the last window on the left)
- Main Status Menu (the last window on the right)

To scroll the information in a status window, or to execute a context-specific DO command, you must make the status window active by pressing the TAB key until that window is highlighted by a thick border. The TAB key moves the active window in sequence from left to right, top to bottom, and then returns to the Edit menu.

Some of the status windows contain more information than can be displayed in the small window. If a lowercase `v` appears in the lower-right corner of a window, it means there is more information available. To scroll through additional information in a window, use the TAB key to move to that window.

See also *Ethernet Status window*, *Line Status windows*, *Main Status Menu window*, *Session Status window*, *Sys Option window*, *System Status window*, *WAN Status window*.

S/T interface—n. The electrical interface between a network terminator (NT1) and one or more ISDN communications devices without their own NT1s. See also *NT1*.

S/T-interface—adj. Describes an ISDN communications device that connects to an external network terminator (NT1). See also *NT1*.

stop bit—In asynchronous transmission, a bit that marks the end of the character. It appears after the parity bit, if one is in use. Compare with *parity bit, start bit*. See also *asynchronous transmission*.

Stop record—A RADIUS-accounting or call-logging record that contains information about the end of a session with the MAX. See also *Stop session*.

Stop session—An event denoting the end of a session with the MAX. Information about the Stop session event appears in a RADIUS-accounting or call-logging Stop record.

STP cable—Shielded Twisted Pair cable. STP cable consists of two wires twisted two or more times per inch in order to help cancel out noise. The entire cable has a protective covering. STP cable is typically used in ARCnet and Token Ring networks. See also *ARCnet*, *Token Ring*

straight-through cable—A cable with wires that have terminating ends with the same wire assignments. Compare with *crossover cable*.

stub area—An Open Shortest Path First (OSPF) area in which all external routes are summarized by a default route. To reduce the cost of routing, OSPF supports stub areas. A stub area allows no Type-5 LSAs to be propagated in the area. Instead, it depends on default routing to external destinations. Compare with *normal area, NSSA*. See also *area, Open Shortest Path First*.

subaddress—A number used for routing incoming calls to the appropriate destination in the MAX unit.

subnet—See *IP subnet*.

subnet mask—An IP feature in which a group of bits identifies a subnet. To specify a subnet mask, the MAX appends to the IP address a modifier that specifies the total number of network bits in the address.

For example, in the address 198.5.248.40/29, the /29 specification indicates that 29 bits of the address specify the network. The three remaining bits specify unique hosts. With three bits used to specify hosts on a 29-bit subnet, eight different bit-combinations are possible:

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

Following are standard and Ascend subnet formats for a class C network number:

Subnet mask	Number of host addresses	Ascend notation
255.255.255.0	254 hosts + 1 broadcast, 1 network base	/24
255.255.255.128	126 hosts + 1 broadcast, 1 network base	/25
255.255.255.192	62 hosts + 1 broadcast, 1 network base	/26
255.255.255.224	30 hosts + 1 broadcast, 1 network base	/27
255.255.255.240	14 hosts + 1 broadcast, 1 network base	/28
255.255.255.248	6 hosts + 1 broadcast, 1 network base	/29
255.255.255.252	2 hosts + 1 broadcast, 1 network base	/30
255.255.255.254	Invalid subnet mask (no hosts)	/31
255.255.255.255	1 host (a host route)	/32

The broadcast address of any subnet has the host portion of the IP address set to all ones. The network address (or base address) represents the network itself, because the host portion of the IP address is all zeros. For example, suppose the MAX configuration assigns the following address to a remote router:

```
198.5.248.120/29
```

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

```
198.5.248.120–198.5.248.127
```

A host route is a special-case IP address with a subnet mask of /32. For example:

```
198.5.248.40/32
```

Host routes are required for a dial-in host.

See also *host number*, *host route*, *IP*, *IP address*, *IP subnet*, *IP network number*.

SubNetwork Access Protocol—See *SNAP*.

subprofile—A set of parameters nested below a top-level profile. See also *CLID*, *profile*.

SuperDigital 128—A nailed-up service available only in Japan. Subscribers receive two ISDN B channels combined into a single 128K pipe.

Supervisory frame—On an X.25 network, a frame that can request and suspend transmission, report on link status, and acknowledge I-frames. Compare with *general frame*, *I-frame*. See also *X.25*.

SVC—Switched Virtual Circuit. An SVC is a path over a packet-switched network. It appears to be a dedicated circuit, but the connection stays up only as long as needed. Compare with *PVC*.

SWIPE—IP with Encryption. SWIPE is a Network-layer security protocol that works by adding a cryptographic authenticator to each packet, and encrypting the data.

switch—A device that connects the calling party to the answering party.

Switched-56—A data service consisting of a single 56-Kbps channel. The Switched-56 data service is available over any type of line. Because Switched-56 was the first available data service, both the service itself and the lines that accessed it were called Switched-56. However, any type of line can now access Switched-56 data service, and there are other new services in addition to Switched-56.

Switched-56 line—A line that provides a single 56-Kbps data channel with inband signaling. See also *inband signaling*.

Switched-64—A data service consisting of a single 64-Kbps channel. The Switched-64 data service is available over T1 PRI and ISDN BRI lines only. See also *ISDN BRI line*, *T1 PRI line.*

Switched-384—A data service consisting of a single 384-Kbps circuit, called an *H0 channel*. The H0 channel is comprised of 6 B channels. The Switched-384 data service is available over T1 PRI lines only. Switched-384 is also known as the *H0 data service*. See also *B channel*, *T1 PRI line.*

Switched-1536—A data service consisting of a single 1536-Kbps circuit, called an *H11 channel*. The H11 channel is comprised of all 24 channels on the line. You must use two T1 PRI lines to access Switched-1536. One line carries the user data, and the other line contains the D channel. Non-Facility Associated Signaling (NFAS) is required for the Switched-1536 data service because the D channel must be on a separate line. The Switched-1536 data service is available over T1 PRI lines only. Switched-1536 is also known as the *H11 data service*. See also *D channel*, *NFAS*, *T1 PRI line*.

switched channel—A channel that provides a temporary connection for the exchange of data. The channel is cleared when the call ends. Compare with *nailed-up channel.*

switched circuit—A temporary connection between endpoints, established for the duration of a call, over which two parties exchange data. The circuit is disconnected when the call ends. Compare with *nailed-up circuit.*

switched line—A line consisting of channels in use only for the duration of the connection. Compare with *nailed-up line*.

Switched Multimegabit Data Service—See *SMDS*.

Switched Nx64—See *MultiRate*.

symbolic name—A name that denotes an IP address. A symbolic name consists of a user name and a domain name in the format *username@domain_name*. The user name corresponds to the host number in the IP address. The domain name corresponds to the network number in the IP address. A symbolic name might be steve@abc.com or joanne@xyz.edu. See also *IP address*.

synchronization—A method of ensuring that the receiving end can recognize characters in the order in which the transmitting end sent them, and can know where one character ends and the next begins. Without synchronization, the receiving end would perceive data simply as a series of binary digits with no relation to one another.

synchronous PPP—A PPP connection that uses synchronous transmission. In Figure 42, a synchronous PPP session takes place between a Pipeline unit and a MAX.

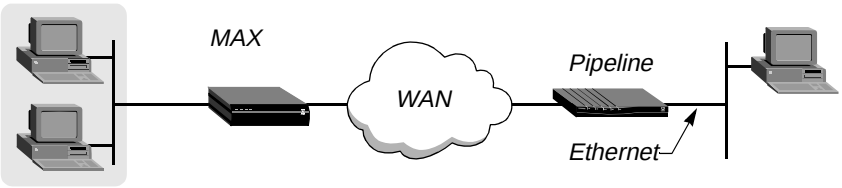

Figure 42. Synchronous PPP connection

Synchronous PPP is also known as *sync PPP*. Compare with *asynchronous PPP*. See also *PPP*, *synchronous transmission*.

synchronous transmission—A transmission mode in which the data moves in large blocks, called messages or frames. A synchronous WAN link uses High-level Data Link Control (HDLC) encoding and connects to a router for a network-to-network link. The MAX routes a synchronous transmission as a digital call to an HDLC channel, and then to the router software. Each synchronous call uses Point-to-Point Protocol (PPP), Multilink Protocol (MP), Multilink Protocol Plus (MP+), or Frame Relay encapsulation.

In a synchronous transmission, both the sending device and the receiving device must maintain synchronization in order to determine where one block of data ends and the next begins. Each side can transmit a separate synchronizing signal, called a clock. Or, each frame or message can contain synchronization information.

In the latter method, each block of data starts with one or more control characters, usually eight bytes long, called a SYNC. The receiver interprets the SYNC as a signal that it can start accepting data. Synchronous transmission can be up to 20 percent faster than asynchronous transmission.

See also *Frame Relay*, *HDLC*, *MP*, *MP+*, *PPP*, *synchronization*.

sync PPP—See *synchronous PPP*.

SYN-to-SYN timer—In an X.25/T3POS configuration, a timer that applies to opening frames in Local or Binary Local mode. Normally, in order to indicate that an idle link is still alive, the PAD sends SYN signals to the Data Terminal Equipment (DTE) at the interval specified by the SYN-to-SYN timer . However, if the DTE sends a SYN signal to the PAD before the PAD sends one to the DTE, the SYN-to-SYN timer specifies the period of time the PAD expects SYN signals from the DTE. If the PAD does not receive two SYN signals with the interval specified by the SYN-to-SYN timer, it tries to restore the link.

The SYN-to-SYN timer is also known as the T2 timer. See also *Binary Local mode*, *Local mode*, *X.25/T3POS*.

Syslog—A facility that sends system status messages to a host computer, known as the *Syslog host*. The Syslog host saves the system status messages in a Syslog file. For detailed information about the Syslog daemon, see the UNIX man pages on `logger(1)`, `syslog(3)`, `syslog.conf(5)`, and `syslogd(8)`. The Syslog function requires User Datagram Protocol (UDP) port 514.

Syslog host—The station to which the MAX sends system logs.

Sys Option window—A status window that contains management information about the MAX. Following is a sample Sys Option window:

```
|--------------------|
|00-100 Sys Option   |
|>Security Prof: 1  ^|
| Software +5.0A0+   |
| S/N: 5210003      v|
|--------------------|
```

The Sys Option window shows which Security profile is active, the Ascend software version that is running, the unit's serial number (S/N), and information on a variety of hardware or software options. The window also displays a system uptime value, which is updated every few seconds to show the number of days, hours, minutes, and seconds the MAX has been operating. For example:

```
Up: 12:17:18:26
```

When the Sys Option window is active, you can use the arrow keys to scroll down and view the list of system options. For example, you see the software load name, various installed software options (such as Frame Relay, AIM, or BONDING), and the AuthServer and AcctServer options, which specify the IP addresses of the RADIUS (or TACACS) authentication server and the RADIUS accounting server, respectively. See also *status window*.

system-based routing—A form of IP routing in which the entire unit has a single IP address. For systems that have a single backbone connection, system-based routing is the simplest way to configure the MAX. Compare with *interface-based routing*.

System Status window—A window that displays messages relating to the system itself. Following is a sample System Status window:

```
l--------------------|
|00-200 15:10:34     |
|>M31  Line    Ch    |
| LAN session up     |
| slc-lab-236        |
|--------------------|
```

The system message log provides a log of up to 32 of the most recent system events. Use the arrow key to scroll up (previous messages) or down (later ones). The Delete key clears all the messages in the log. The message log window is organized as follows:

- The first line shows the menu number and the time the most recently logged event occurred.
- The second line identifies the log entry number (M00-M31) and, if applicable, the line and channel on which the event occurred.
- The third line contains the text of the message. For example:

  `Call Terminated` means an active call disconnected normally.

  `LAN session up` means that an incoming connection has been established.

  `No Connection` means the remote device did not answer the call.
- The fourth line contains a message qualifier, such as a name or phone number that qualifies the message displayed.

See also *status window*.

T

T1 channel—One of 24 channels on a T1 line. See also *fractional T1 line, T1 line, T1 PRI line, unchannelized service.*

T1 line—A line that supports 24 64-Kbps channels, each of which can transmit and receive data or digitized voice. The line uses framing and signaling to achieve synchronous and reliable transmission. The most common configurations for T1 lines are ISDN Primary Rate Interface (T1 PRI) and unchannelized T1, including fractional T1. The MAX supports up to four T1 lines for up to 96 concurrent sessions. See also *fractional T1 line, T1 channel, T1 PRI line, unchannelized service.*

T1 PRI line—T1 Primary Rate Interface line. A T1 PRI line has a total bandwidth of 1.544 Mbps. It uses 23 B channels for user data, and one 64-Kbps D channel for ISDN D-channel signaling. The B channels can be all switched, all nailed up, or a combination of switched and nailed up. The T1 PRI line is a standard in North America, Japan, and Korea. You can connect this type of line to standard voice, Switched-56, Switched-64, Switched-384, Switched-1536, and MultiRate data services. Using a feature called PRI-to-TI conversion, the MAX can share the bandwidth of a T1 PRI line with a PBX. Compare with *E1 PRI line, ISDN BRI line, unchannelized service*. See also *B channel, D channel, MultiRate, nailed-up channel, PCM, Switched-56, Switched-64, Switched-384, Switched-1536, T1 channel, T1 line*.

T1 Primary Rate Interface line—See *T1 PRI line*.

T1 retransmission timer—In an X.75 configuration, the maximum amount of time in ticks the transmitter should wait for an acknowledgment before initiating a recovery procedure. See also *X.75*.

T3POS—Transaction Processing Protocol for Point-of-Service. T3POS is a character-oriented, frame-formatted protocol designed for Point-of-Service (POS) transactions through an X.25-based packet-switched network. T3POS allows you to send data over the ISDN D channel while continuing to send traffic over both B channels. The T3POS protocol involves three parties—the T3POS DTE , the T3POS PAD, and the T3POS Host (Figure 43).

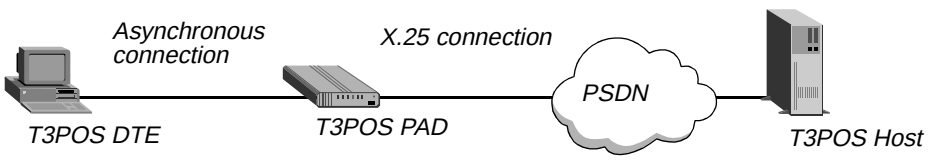

Figure 43. T3POS setup

A typical use of T3POS is to perform credit-card authorization over the D channel while using the B channels to transmit inventory-control data and other traffic. See also *X.25*.

T3 timer—See *ENQ handling timer*.

T4 timer—See *Response timer*.

T5 timer—See *DLE, EOT timer*.

T302 timer—The amount of time that the MAX waits for an Info message from the switch when the Sending Complete Information Element (IE) is not in the Setup message.

The layer 3 Setup message consists of many IEs, such as Bearer Capability IE, Channel Identifier IE, Caller Number IE, Called Number IE, Sending Complete IE, and so on. The MAX checks for the Sending Complete IE upon receiving the Setup message from the switch. If the Sending Complete IE is not in the Setup message, the MAX starts the T302 timer and waits for an Info message from the switch. If the Info message consists of a Sending Complete IE, the MAX stops the T302 timer. If no Sending Complete IE appears, the MAX restarts the T302 timer.

See also *ISDN Call Setup message*.

TA—Terminal Adapter. A TA is a protocol converter that adapts non-ISDN equipment (such as a phone, fax, or modem), and enables each device to work over an ISDN connection. A TA has two functions. First, it must change the format of transmitted data to match either the V.110 or V.120 standard for asynchronous transfer over a B channel. Second, it must provide a way of setting up and clearing calls, usually by means of Hayes AT commands.

A TA is to an ISDN line what a modem is to an analog telephone line. However, some of the D-channel information does not pass through the TA, so non-ISDN equipment cannot take full advantage of ISDN facilities, such as Calling-Line ID (CLID).

See also *ISDN*, *V.120*.

TACACS—Terminal Access Concentrator Access Control Server. TACACS is a very simple query/response protocol that enables the MAX to check a user's password in order to grant or prevent access. A TACACS server supports only the basic password exchanges that Password Authentication Protocol (PAP) uses. It does not support Challenge Handshake Authentication Protocol (CHAP). See also *CHAP*, *PAP*.

TACACS+—Terminal Access Concentrator Access Control Server Plus. TACACS+ is a proprietary Cisco enhancement to the Terminal Access Concentrator Access Control Server (TACACS) protocol. TACACS+ handles the transfer of authentication and authorization information between a Network Access Server (NAS) and an authentication server, encrypting password information and forwarding it over the network. TACACS+ supports AppleTalk Remote Access (ARA), Challenge Handshake Authentication Protocol (CHAP), Password Authentication Protocol (PAP), Point-to-Point Protocol (PPP), Serial Line Internet Protocol (SLIP), and Telnet. In addition, TACACS+ uses the TCP protocol to transmit accounting information to an accounting server. See also *CHAP*, *PAP*, *PPP*, *SLIP*, *TCP*, *Telnet*.

tag—An Open Shortest Path First (OSPF) method of flagging a route as external—that is, as having been imported into the OSPF database from outside the router's Autonomous System (AS). See also *AS*, *external route*, *OSPF*.

TAOS—True Access Operating System. TAOS is Ascend's comprehensive software architecture for WAN access. It consists of two main components: the standard TAOS kernel and optional TAOS extensions.

The TAOS kernel comes standard on all MAX WAN-access switches, and supports IP routing, modem management, terminal-server functionality, authentication, authorization, accounting, WAN protocols, and bandwidth management.

TAOS extensions allow service providers and corporations to further customize and enhance their WAN-access services. The extensions include the following:

- Global Digital Access for ISDN and Frame Relay environments, including support for ISDN clients, Frame Relay concentration, and ISDN signaling.
- IntragyCentral, the embedded WAN-access switch component of Ascend's Intragy enterprise access software suite. IntragyCentral provides multiprotocol routing, multiprotocol dial access, transparent bridging, and modem pooling for LAN-based outbound dial and fax.
- SecureConnect, Ascend's state-of-the-art security technology, providing an integrated, dynamic firewall and eliminating the need for separate security hardware.
- Manageability under NavisAccess, the leading management platform for access environments.
- Tunneling support for Virtual Private Networks (VPNs).
- Quality of Service (QoS) solutions, including videoconferencing support.

tariff—A document filed by a regulated telephone company with a state public utility commission or the Federal Communications Commission. A tariff details services, equipment, and pricing publicly offered by the telephone company.

TCP—Transmission Control Protocol. TCP operates at the Transport layer, providing connected-oriented services. It uses IP to deliver packets. See also *IP*.

TCP-Clear—See *Raw TCP*.

TCP/IP—Transmission Control Protocol/Internet Protocol. TCP/IP is a family of protocols that defines the format of data packets sent across a network, and is the communications standard for data transmission between different platforms.

The TCP/IP family consists of the following protocols and services:

OSI layer	Protocol name	Description of service
Application	Boot Protocol (BOOTP) File Transfer Protocol (FTP) Telnet	User services that provide applications a computer can use
	Network File System (NFS) Network Information Service (NIS) Remote Procedure Call (RPC) Simple Mail Transfer Protocol (SMTP) Simple Network Management Protocol (SNMP)	File-transfer, mail, and management services
Session	Border Gateway Protocol version 4 (BGP) Gateway-to-Gateway Protocol (GGP) Interior Gateway Protocol (IGP)	Gateway protocols that enable networks to share routing and status information
Transport	Transmission Control Protocol (TCP) User Datagram Protocol (UDP)	Transport protocols that control data transmission between computers
Network	Internet Protocol (IP) Internet Control Message Protocol (ICMP) Routing Information Protocol (RIP) Open Shortest Path First (OSPF)	Routing protocols that control addressing and packet assembly, and determine the best route for a packet to take to arrive at its destination
	Domain Name System (DNS) Address Resolution Protocol (ARP) Reverse Address Resolution Protocol (RARP)	Network-address services and protocols that handle the way each computer on a network is identified

See also *Address Resolution Protocol*, *BGP*, *BOOTP*, *DNS*, *EGP*, *FTP*, *GGP*, *ICMP*, *IGP*, *IP*, *NFS*, *NIS*, *OSPF*, *RARP*, *RIP*, *RPC*, *SMTP*, *SNMP*, *TCP*, *Telnet*, *UDP*.

TCP/IP header compression—See *VJ compression*.

TCP timeout—The length of time the MAX attempts to connect to an IP host in the list provided by the DNS server. Because the first host on the list may not be available, the timeout should be short enough to allow the MAX to go on to the next address on the list before the client software times out. This feature applies to all TCP connections initiated from the MAX, including Telnet, Rlogin, TCP-Clear, and the TCP portion of DNS queries. See also *DNS*, *Raw TCP*, *Rlogin*, *Telnet*.

TD—Transmit Data. TD is a signal that indicates that the modem is sending data. See also *DB25 pin connector*

TDM—Time Division Multiplexing. TDM is a scheme that uses time-slot assignment, enabling information from multiple channels to use bandwidth on a single line.

TE—Terminal Equipment. A TE is any ISDN-compatible device attached to a network, such as a PBX, telephone, fax machine, or computer.

Telecommunications Industry Association—See *TIA*.

telecommuter—A work-at-home computer user who connects to the corporate LAN backbone by means of remote-access technology. For example, a telecommuter can establish a link with the LAN by means of a modem connected to an analog line, an ISDN Terminal Adapter (TA) or router connected to an ISDN line, or a Channel Service Unit/Data Service Unit (CSU/DSU) connected to a Switched-56 line. See also *analog line*, *CSU*, *DSU*, *ISDN line*, *modem*, *TA*.

Telnet—A protocol that links two computers in order to provide a terminal connection to the remote machine. Instead of dialing into the computer, you connect to it over the Internet using Telnet. When you issue a Telnet session, you connect to the Telnet host and log in. The connection enables you to work with the remote machine as though you were a terminal connected to it. You can remotely manage your MAX by establishing a Telnet session to the MAX from any Telnet workstation on the network. See also *Telnet host*, *Telnet session*.

Telnet host—A device with which you establish a Telnet session. See also *Telnet*, *Telnet session*.

Telnet mode—The mode in which terminal-server Telnet users communicate with the Telnet host. See also *ASCII mode*, *Binary mode*, *Transparent mode.*

Telnet password—The password a user must enter to access the MAX unit by means of Telnet. A user is allowed three tries of sixty seconds each to enter the correct password. See also *Telnet*, *Telnet host*, *Telnet session*.

Telnet session—A terminal connection to a remote machine by means of the Telnet protocol. After you set up a basic IP configuration for the MAX, users can Telnet into the MAX. Each user can initiate a Telnet session to the MAX from a local workstation or from a WAN connection. In both cases, the MAX authenticates the session. In addition to the MAX password, you can specify that Telnet requires its own password authentication, which occurs prior to any MAX authentication. See also *Telnet*, *Telnet host*.

terminal—A computer that does not have its own processor and that must connect to a terminal server in asynchronous mode to use its Central Processing Unit (CPU). VT100, ANSI, and TTY are all types of terminals.

Terminal Adapter—See *TA*.

Terminal Access Concentrator Access Control Server—See *TACACS*.

Terminal Access Concentrator Access Control Server Plus—See *TACACS+*.

terminal emulation program—See *terminal emulator*.

terminal emulator—A program that makes your computer look like a terminal so that you can connect to a terminal server. Your computer acts like a terminal during the connection. All processing is taking place remotely. A terminal emulator is also called a *terminal emulation program*.

Terminal Equipment—See *TE*.

terminal mode—A terminal-server access mode in which the MAX negotiates a user-to-host session. Instead of providing only the login name and password required to authenticate a Connection profile or RADIUS user profile, you can set up an expect-send script that also includes the terminal-server prompt and a command, such as PPP, SLIP, TCP, Telnet, or Rlogin. In this way, the session with a host comes about as part of the login process, so the user never actually sees the terminal-server command-line prompt. Alternatively, you can provide access to the command line and restrict the commands you make accessible to the user. See also *expect-send script*, *PPP*, *Rlogin*, *SLIP*, *TCP*, *Telnet*, *terminal server*.

terminal server—A terminal server is a computing device to which a terminal can connect over a LAN or WAN link. A terminal communicates with the terminal server over an asynchronous serial port (typically an RS-232 port) through a modem. A terminal converts the data it receives from the terminal server into a display and does no further processing of the data. A terminal also converts the operator's keystrokes into data for transmission to the terminal server.

The MAX terminal-server software receives asynchronous calls after they have been processed by a digital modem. Typically, a modem or V.120 Terminal Adapter (TA) dials these calls. V.120 and TCP calls are enabled by default. If the caller does not send Point-to-Point Protocol (PPP) packets immediately, the terminal server starts a login sequence.

Figure 44 shows an incoming modem call. A PC running SoftComm initiates the connection. (SoftComm is a program that causes the user's modem to dial into the MAX.) The MAX directs the call to its digital modem, and then forwards the calls to its terminal-server software. In Figure 44, the MAX immediately directs the call to a Telnet host.

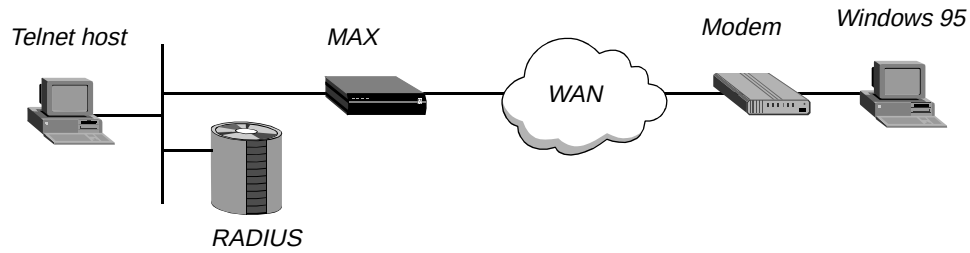

Figure 44. Terminal-server connection

Each user must have a Connection profile or a RADIUS user profile that specifies a name and password to use in the terminal-server login sequence. When it receives a name and password from the caller, the terminal server authenticates them by means of a Connection profile or external authentication server, and then performs one of the following actions:

- Displays the terminal-server command-line prompt
- Displays a menu of hosts the user can log into
- Immediately logs the user into a designated host
- Initiates a PPP or SLIP session with the user

To protect the command-line from unauthorized access, you can also choose to assign the terminal server its own password.

If it receives an asynchronous PPP call, the terminal server does not begin a login sequence. Instead, it responds with a PPP packet, and Link Control Protocol (LCP) negotiation begins, including negotiation for Password Authentication Protocol (PAP) or Challenge Handshake Authentication Protocol (CHAP) authentication. The terminal server directs the call to the router software, and the connection proceeds as for a regular synchronous PPP session. The user bypasses the terminal server interface altogether.

In most cases, the terminal server is a stepping stone toward access to one or more network hosts. To enable host access, you can configure the terminal server in terminal mode, immediate mode, or menu mode.

See also *asynchronous PPP*, *CHAP*, *Connection profile*, *digital modem*, *immediate mode*, *menu mode*, *modem*, *PAP*, *PPP*, *SLIP*, *terminal mode*, *user profile*, *V.120 TA*.

terminal-server connection—A connection between a terminal and a terminal server over a LAN or WAN link. See also *terminal server*.

terminal-server idle timer—A specification that determines the number of seconds a terminal-server connection must remain idle before the MAX disconnects the session. See also *terminal server*, *terminal-server session*.

terminal-server session—An end-to-end connection between a terminal and a terminal server. Usually, the terminal-server session begins when the call goes online and ends when the call disconnects. The MAX supports all the common capabilities of standard terminal servers, including Telnet, Domain Name System (DNS), login and password control, Call Detail Reporting (CDR), and authentication services. See also *terminal server*.

Terminal Timing signal— Specified in the V.35, X.21, and RS-449 serial interfaces, a clock signal that compensates for the phase difference between Send Data and Send Timing. For the MAX to use the Terminal Timing signal from the codec, the AIM port module must support Terminal Timing and the codec must use Terminal Timing if the distance and serial data rate between the MAX and the host is greater than the following:

- A maximum cable length of 25 feet and a serial data rate of 3 Mbps
- A maximum cable length of 75 feet and a serial data rate of 2 Mbps
- A maximum cable length of 150 feet and a serial data rate of 512 Kbps

See also *RS-449*, *V.35*, *X.21*.

TFTP—Trivial File Transfer Protocol. TFTP is a simplified version of FTP. It enables you to transfer files from one computer to another.

Thick Ethernet—A type of .4" diameter coaxial cable for Ethernet networks. Also known as *thicknet*.

thicknet—See *Thick Ethernet*.

Thin Ethernet—A type of .2" diameter coaxial cable for Ethernet networks. Also known as *thinnet*.

thin load—Version 4.6Ci19 to 5.0Aix and all 5.0Ap releases with a file size less than 960 KB (for MAX units) or 448 KB (for Pipeline units). Compare with *fat load*.

thinnet—See *Thin Ethernet*.

third-party routing—A feature that enables the MAX to advertise Open Shortest Path First (OSPF) routes to external destinations on behalf of another gateway, commonly known as advertising a forwarding address. When third-party routing is enabled, the MAX advertises the IP address of another gateway. If third-party routing is disabled, the MAX advertises itself as the forwarding address to an external destination.

Depending on the exact topology of the network, other routers might be able to route packets directly to the forwarding address without involving the advertising MAX, increasing the total network throughput. In this scenario, all OSPF routers must know how to route to the forwarding address.

See also *OSPF*.

TIA—Telecommunications Industry Association. The TIA is a group that determines standards for the electrical level of data transmission.

tick—An IBM unit of measurement that corresponds to one-eighteenth of a second.

Time Division Multiplexing—See *TDM*.

timeout—An event in which a device or user exceeded a configured time limit for responding to a device or process.

token—A password that appears in the LCD display of a token card. See also *token card*, *token-card authentication*, *token-card server*.

token card—A hardware device, typically shaped liked a credit-card calculator, that displays a current, one-time-only password (called a *token*). The token grants a user access to a secure network, and changes many times per day. Token cards keep changing authentication information continuously up-to-date by maintaining a synchronized clock with a token-card server, such as a Security Dynamics ACE/Server, an AssureNet Pathways Defender server, or an Enigma Logic SafeWord server. To gain access to a secure network, each authorized user must have a token card.

A token card protects against replay attacks, in which an unauthorized user records valid authentication information exchanged between systems, and then replays it later to gain entry. Because the token is a one-time-only password, replay is impossible. See also *ACE authentication*, *Defender authentication*, *SafeWord authentication*, *token*, *token-card authentication*, *token-card server*.

token-card authentication—A form of authentication requiring that users change passwords many times per day. When connecting to a Security Dynamics ACE/Server or an Enigma Logic SafeWord server, the MAX supports token-card authentication by using a RADIUS server as the intermediary between the MAX and the authentication server. When communicating with an AssureNet Pathways Defender server, the MAX does not use RADIUS as an intermediary. Rather, the MAX communicates directly with the Defender server.

Figure 45 shows a dial-in connection to the MAX. The remote user must use a token card to gain access to the secure network.

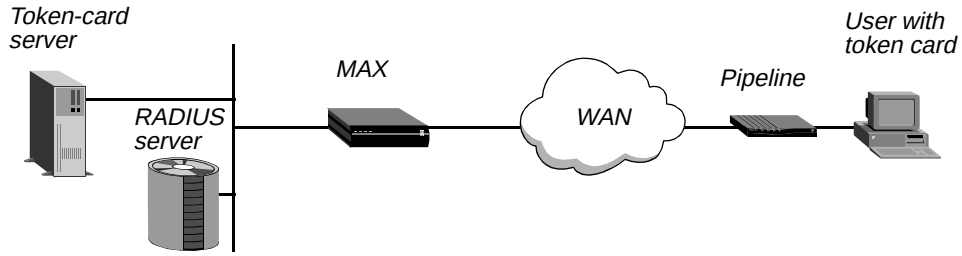

Figure 45. Token-card authentication for dial-in connections

Figure 46 shows a dial-out connection from the MAX. The local user must use a token card to gain access to the remote secure network.

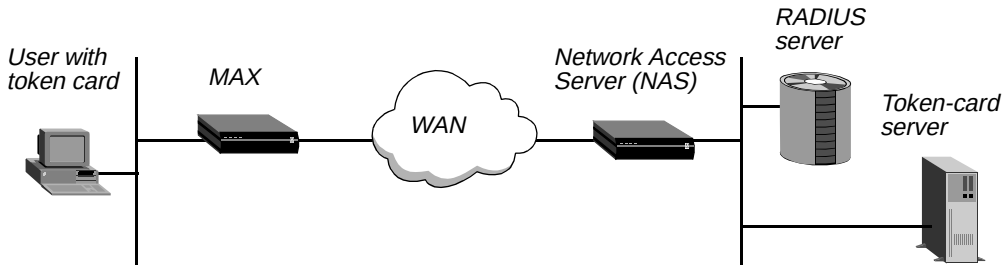

Figure 46. Token-card authentication for dial-out connections

A local user with a token card initiates a connection by logging into the MAX terminal server and dialing out to the remote unit. These actions require that the user have login privileges to the MAX unit, and that the MAX have a Connection profile or RADIUS profile configured for a connection to the remote device.

See also *ACE authentication*, *Defender authentication*, *SafeWord authentication*, *token*, *token card*, *token-card server*.

token-card server—A server that maintains a synchronized clock with hand-held token cards to provide users with a current, one-time-only password (called a *token*). The correct token is required to gain access to a secure network. Examples of token-card servers are the Security Dynamics ACE/Server, the AssureNet Pathways Defender server, and the Enigma Logic SafeWord server. See also *ACE authentication*, *Defender authentication*, *SafeWord authentication*, *token*, *token card*, *token-card authentication*.

Token Ring—A network architecture that uses a ring topology, baseband signaling, and the token-passing media-access method. Token Ring can operate at 1, 4, or 16 Mbps, and supports four-wire twisted pair or fiber-optic media.

TOS—Type-of-Service. A feature that enables an Internet device to select the quality of service. The TOS is specified along the abstract values of precedence, delay, throughput, reliability, and cost.

You can configure the MAX to set priority bits and TOS classes of service on behalf of customer applications. The MAX does not implement priority queuing, but it does set information that can be used by upstream routers to prioritize and select links for particular data streams. Specifically, you can define a policy in a Connection or RADIUS profile. The parameters in the profile set bits in the TOS byte of IP packet headers that are received on the WAN interface, transmitted on the WAN interface, or both. Another router can then interpret the bits accordingly. You can also define a TOS filter, which can then be applied to any number of Connection or RADIUS profiles. .

For a Connection or RADIUS profile that has both its own local policy and an applied TOS filter, the policy defined in the TOS filter takes precedence. For example, applying a TOS filter to a TOS-enabled connection allows you to define one priority setting for incoming packets on a connection and another for incoming packets addressed to a particular destination (the destination in a TOS filter).

See also *precedence*, *TOS filter*.

TOS filter—Type-of-Service filter. A TOS filter is a packet filter that enables you to specify many of the same values as an IP filter, in addition to a precedence and TOS value. Like other kinds of Ascend packet filters, a TOS filter can affect incoming packets, outgoing packets, or both. Compare with *IP filter*. See also *packet filter*, *precedence*.

Transaction Procession Protocol for Point-of-Service—See *T3POS*.

translation table—A table used by Network Address Translation (NAT) for LAN. The translation table is limited to 500 addresses. A translation table entry represents one TCP or UDP connection. The translation table entries are reused as long as packets are seen that match an entry. All are freed (expired) when a connection disconnects.

Transmission Control Protocol—See *TCP*.

Transmission Control Protocol/Internet Protocol—See *TCP/IP*.

Transmit Data—See *TD*.

transmit data rate—The baud rate of data sent by the MAX. Compare with *receive data rate*.

transparent bridge—A bridge that notes a packet's source address and creates a bridging table that associates a host's Media Access Control (MAC) address with a particular Ethernet interface. The MAX is an example of a transparent bridge (also called a *learning bridge*). See also *bridge*, *bridging*.

Transparent mode—A data-transfer mode for host-initiated calls on an X.25/T3POS network; also, a Telnet mode. In Transparent mode for an X.25/T3POS network, the T3POS PAD does not provide any error recovery. Rather, the Data Terminal Equipment (DTE) and the host system provide error recovery for the connection. In Transparent mode, however, the T3POS PAD does recognize a DLE, EOT command from the DTE, and clears the call when it receives one.

In Transparent mode for a Telnet connection, you can send and receive binary files without having to be in Binary mode. You can run the same file transfer protocols available in Binary mode.

Compare with *ASCII mode*, *Binary mode*, *Binary Local mode*, *Blind mode*, *Local mode*. See also *DLE*, *EOT command*, *DTE*, *PAD*, *X.25/T3POS*.

Transport layer—The middle layer of the OSI Reference Model. The Transport layer provides data transfer at the proper speed, quality, and error rate, ensuring reliable delivery. Examples of Transport-layer protocols are Transmission Control Protocol (TCP) and User Datagram Protocol (UDP). See also *OSI Reference Model*, *TCP*, *UDP*.

traps-PDU—A message that a Simple Network Management Protocol (SNMP) agent sends to a manager application to inform the manager of network events. See also *agent*, *community name*, *manager*, *MIB*, *SNMP*.

Trivial File Transfer Protocol—See *TFTP*.

True Access Operating System—See *TAOS*.

trunk group—A group of switched channels to use for outgoing calls. To specify that outgoing calls use a specific bandwidth, you can configure a Connection profile or Call-Route profile to refer to a specific trunk group. You can also use trunk groups to separate lines supplied by different carriers. Each set of lines can be used as a backup if a switch becomes unavailable.

The decision to use trunk groups is a global one. Once you have enabled the use of trunk groups, every switched channel must be assigned a trunk group number or it will not be available for outgoing calls. In addition, trunk groups limit the number of channels available to multichannel calls, because only channels within the same trunk group can be aggregated.

trunk-side connection—A line that extends from the telephone company's Central Office (CO) to the telephone network. Typically, a trunk-side connection is high-bandwidth and all digital. Compare with *line-side connection.*

tunneling—A way of overcoming protocol restrictions on a network by encapsulating packets that use an unsupported protocol inside packets that use a protocol supported by the network.

Tunnel Server—When a tunneling protocol is in use, the system that decapsulates the packets. Examples of tunneling protocols are Ascend Tunnel Management Protocol (ATMP), Layer 2 Tunneling Protocol (L2TP), and Point-to-Point Tunneling Protocol (PPTP). See also *ATMP*, *L2TP*, *PPTP*.

twisted-pair cable—A cable consisting of four or more copper wires twisted together in pairs. Telephone wiring is an example of twisted-pair cable. Twisted-pair cable can be shielded or unshielded. See also *STP cable*, *UTP cable.*

twisted-pair Ethernet—See *10Base-T*.

two-wire subscriber loop—A two-wire WAN link connecting the Customer Premises Equipment (CPE) to the carrier's switch. See also *CPE*.

Type-5 LSA—See *ASE Type-5*.

Type-7 LSA—See *ASE Type-7*.

Type-of-Service—See *TOS*.

U

UART—Universal Asynchronous Receiver/Transmitter. A UART is a chip that provides a RS-232C Data Terminal Equipment (DTE) interface to a device, enabling the unit to communicate with its attached serial devices. See also *asynchronous transmission, DTE, RS-232C, serial transmission*.

UDP—User Datagram Protocol. UDP is a Transport-layer protocol that provides connectionless service without packet acknowledgment. See also *Transport layer, UDP port.*

UDP port—A16-bit number that allows multiple processes to use User Datagram Protocol (UDP) services on the same host. A UDP address is the combination of a 32-bit IP address and the 16-bit port number. Examples of well-known UDP ports are 7 (for Echo packets), 161 (for SNMP packets), and 514 (for Syslog packets). See also *UDP*.

UDP queue—A queue containing unprocessed User Datagram Protocol (UDP) requests. Compare with *backoff queue, RIP queue, SNMP queue*. See also *queue*.

U interface—n. The electrical interface between an ISDN telephone line and a network terminator (NT1) device. See also *NT1*.

U-interface—adj. Describes an ISDN communications device that connects directly to an ISDN telephone line. A U-interface device contains its own network terminator (NT1). See also *NT1*.

U-Law—An ITU-T standard for sampling data by means of Pulse Coded Modulation (PCM). U-Law is most commonly used in North America and Japan. U-Law is also known as *Mu-Law.* Compare with *A-Law*. See also *PCM*.

unchannelized E1—See *unchannelized service*.

unchannelized service—A service that uses the entire bandwidth of a T1 PRI line (1.544 Mbps) or an E1 PRI line (2.048 Mbps). You can use an unchannelized line for a nailed-up connection, such as the link to a Frame Relay network. The MAX treats the line as though it were a single connection at a fixed speed, without individual channels. See also *E1 PRI line, T1 PRI line*.

unchannelized T1—See *unchannelized service*.

UNI—User-to-Network Interface. UNI is the interface between an end user and a network endpoint (a router or a switch) on the Frame Relay network. Data Terminal Equipment (DTE) and Data Circuit-Terminating Equipment (DCE) devices perform UNI procedures for link management. These procedures occur in one direction—the DTE requests information and the DCE provides it. When it is configured with a UNI to Frame Relay, the MAX can act as the user side (UNI-DTE) communicating with the network side (UNI-DCE) of a switch., or as the network side (UNI-DCE) communicating with the user side (UNI-DTE) of a Frame Relay device. Compare with *NNI*. See also *DCE, DTE, Frame Relay network, UNI-DCE interface, UNI-DTE interface*.

unicast network—A network in which a router sends packets to one user at a time. Compare with *broadcast network, multicast network.*

UNI-DCE interface—User-to-Network Interface—Data-Circuit-terminating-Equipment Interface. On a UNI-DCE interface, the MAX acts as the network side communicating with the user side (UNI-DTE) of a Frame Relay device. Figure 47 shows an example of the MAX with a DCE interface.

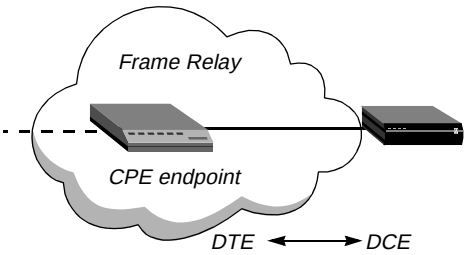

Figure 47. Frame Relay DCE interface

Compare with *UNI-DTE interface*. See also *DCE, DLCI, DTE, Frame Relay, Frame Relay network, UNI.*

UNI-DTE interface—User-to-Network Interface—Data-Terminal-Equipment Interface. On a UNI-DTE interface, the MAX acts as the user side communicating with the network side DCE switch. It initiates link-management functions by sending a Status Enquiry to the UNI-DCE device. Status Enquiries may include queries about the status of Permanent Virtual Connection (PVC) segments the DTE knows about, as well as the integrity of the datalink between the UNI-DTE and UNI-DCE interfaces.

Figure 48 shows an example of the MAX with a UNI-DTE interface.

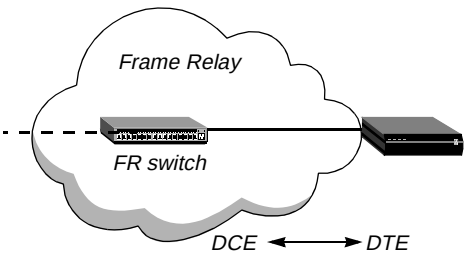

Figure 48. Frame Relay DTE interface

Compare with *UNI-DCE interface*. See also *DCE, DTE, Frame Relay switch, UNI.*

Universal Asynchronous Receiver/Transmitter—See *UART*.

UNIX—A 32-bit operating system that allows multiple users to share resources and perform multiple tasks at the same time. UNIX was developed at Bell Laboratories in 1969. Its development has occurred along two lines: the AT&T System versions and the UC Berkeley Distribution (BSD) releases. The two strains were combined by the UNIX Systems Group into System V Release 4.2 (SVR 4.2).

UNIX DBM database—See *DBM database*.

UNIX hosts file—The `/etc/hosts` file on the UNIX host. The UNIX hosts file contains the names and IP address of all the hosts with which the UNIX server can communicate.

UNIX password—A password entered in the `/etc/password` file on the UNIX host. In a RADIUS user profile, setting the password to `UNIX` provides authentication through the normal UNIX authentication procedure. You cannot specify a UNIX password with Challenge Handshake Authentication Protocol (CHAP) authentication. See also *CHAP*.

UNIX password file—The `/etc/password` file on the UNIX host. The UNIX password file contains passwords for standard UNIX authentication.

unnumbered interface—A link that uses system-based routing, in which the entire MAX system has a single IP address. If all interfaces are unnumbered, the MAX operates as a purely system-based router. Compare with *interface-based routing*, *numbered interface*. See also *IP routing*, *system-based routing*.

Unshielded Twisted Pair cable—See *UTP cable*.

upstream path—The path a call takes from the end user's home to the carrier's Central Office (CO).

User Datagram Protocol—See *UDP*.

user facility—See *facility*.

user name—The name a user must enter to access the services of the MAX. See also *password*.

user profile—A RADIUS users file entry that contains authentication, incoming call configuration, dialout, routing, and filter information. Each user profile consists of a series of attributes. The attributes indicate a user name and password, and enable you to configure routing, bridging, call management, and restrictions on the types of MAX resources a caller can access. See also *pseudo-user profile*, *RADIUS*, *RADIUS server*, *users file*.

users file—A RADIUS file that contains a set of user and pseudo-user profiles. A user profile enables RADIUS to authenticate a dial-in user. It consists of attributes describing the user, and the services he or she can access. A pseudo-user profile is an entry containing information that the MAX can query. It does not exist for the purpose of authenticating a user. Rather, it enables you to specify static route configurations, Frame Relay profile information, bridging entries, and other settings.

A users file can have a flat ASCII format or a UNIX DBM database format. RADIUS must search a flat ASCII file sequentially, which might increase access time, especially if you have many users and many authentication requests. If you use the DBM database version of the users file, RADIUS can locate a record by index with only a few database accesses. However, if you reset passwords, the new passwords take effect only after you rebuild the database. If resetting expired passwords is an important component of your system, the flat ASCII file might be the better choice.

See also *DBM database*, *flat ASCII users file*, *pseudo-user profile*, *RADIUS*, *RADIUS server*, *user profile*.

User-to-Network Interface—See *UNI*.

UTC—Coordinated Universal Time. Formerly known as Greenwich Mean Time (GMT), UTC is the time at the Greenwich observatory, used as a reference point for calculating standard time values.

UTP cable—Unshielded Twisted Pair cable. UTP cable consists of two wires twisted two or more times per inch in order to help cancel out noise. The entire cable has no covering. UTP cable is typically used in telephone lines for voice service, ARCnet networks, 10Base-T Ethernet networks, and particular sections of Token Ring networks. See also*10Base-T*, *ARCnet*, *Token Ring*.

UTP Ethernet—See *10Base-T*.

V

V.21—An ITU-T communications standard for 300-bps full-duplex modems.

V.22—An ITU-T communications standard that supports a data rate of up to 1200 bps at 600 baud.

V.22bis—An extension of the V.22 standard, providing a data rate of up to 2400 bps at 600 baud. See also *V.22*.

V.23—An ITU-T communications standard for 600-bps and 1200-bps full-duplex modems.

V.24—An ITU-T standard that specifies a Physical-layer interface between Data Terminal Equipment (DTE) and Data Circuit-terminating Equipment (DCE). V.24 is nearly identical to RS-232. See also *DCE*, *DTE*, *RS-232*.

V.25bis—An ITU-T communications standard for automatic calling and answering equipment on the Public Switched Telephone Network (PSTN).

V.32—An ITU-T communications standard for full-duplex modem transmission of data across phone lines at rates of up to 9600 bps, with a fallback rate of 4800 bps. A V.32 modem automatically adjusts its transmission speed based on the quality of the line. Compare with *V.34*.

V.32bis—An extension of the V.32 standard, providing a data rate of up to 14,400 bps or fallback to 12,000, 9600, 7200, and 4800 bps. See also *V.32*.

V.34—An ITU-T communications standard for full-duplex modem transmission of data across phone lines at rates of up to 28,800 bps. A V.34 modem automatically adjusts its transmission speed based on the quality of the line. Compare with *V.32*.

V.34bis—An extension of the V.34 standard, providing a data rate of up to 33,600 bps. See also *V.34*.

V.35—An ITU-T standard for high-speed synchronous data transmission and exchange. In the U.S., most routers and Data Service Units (DSUs) that connect to T1 lines use V.35. See also *DSU*, *router*, *synchronous transmission*, *T1 line*.

V.42—An ITU-T error-detection standard for high-speed modems over digital telephone lines. The V.42 standard makes use of the Link Access Procedure, Modem (LAPM). See also *LAPM*.

V.42bis—An ITU-T data compression standard for use with V.42 technology. See also *V.42*.

V.90—An ITU-T communications standard that provides up to 56,000 bps downstream. V.90 was developed on the basis of 3Com's x2 technology and Rockwell's K56flex devices.

V.110—A rate-adaption standard, based on fixed frames, that subdivides the ISDN channel so that it can carry one lower-speed data channel.

V.110 card—A MAX card that provides up to eight V.110 WAN sessions. You can install a maximum of six V.110 cards in the MAX. See also *V.110*.

V.110 TA—V.110 Terminal Adapter. A V.110 TA is a device that changes the format of asynchronous data to match the specifications of the V.110 standard for data transmission over an ISDN line. See also *TA, V.110*.

V.120—A standard for encapsulating asynchronous data communication into synchronous ISDN data. Using standard, asynchronous-only COM ports and a V.120 adapter, two computers can communicate over an ISDN connection. The V.120 adaptor can be connected externally or internally.

V.120 TA—V.120 Terminal Adapter. A V.120 TA is an asynchronous device that changes the format of asynchronous data to match the specifications of the V.120 standard for data transmission over an ISDN line. A V.120 TA is also known as an *ISDN modem*. See also *TA, V.120*.

Van Jacobson compression—See *VJ compression*.

Variable-Length Subnet Mask—See *VLSM*.

VC—Virtual Circuit. A VC is a logical, bidirectional data path between two endpoints. See also *PVC, SVC*.

VCE timer—Virtual Call Establishment timer. On an X.25/PAD network, the VCE timer specifies the number of seconds to maintain a connection to a character-oriented device (such as the terminal server) that has not established a virtual call. See also *X.25/PAD*.

Vendor-Specific Attribute support—See *VSA support*.

videoconferencing—The use of a digital video-transmission system to communicate by means of video and voice. A digital video-transmission system typically consists of a camera, codec, network-access equipment, network, and audio system.

virtual AppleTalk network—A network required for the MAX to route AppleTalk to dial-in clients. You define a virtual AppleTalk network by defining a unique network range. See also *AppleTalk routing, network range*.

Virtual Call Establishment timer—See *VCE timer*.

virtual IPX network—A network required for the MAX to route IPX to dial-in clients. When a NetWare client dials in, the MAX negotiates a routing session by assigning the client a network address on the virtual IPX network. The client must accept the network number that the MAX assigns. If the client has its own node number, the MAX uses that number to form the full network:node address. If the client does not have a node number, the MAX assigns it a unique node address on the virtual network. See also *IPX network*.

Virtual Private Network—See *VPN*.

VJ compression—Van Jacobson compression. VJ compression is a method for compressing Transmission Control Protocol (TCP) headers in order to decrease round-trip times on Serial Line Internet Protocol (SLIP) connections. The version of SLIP implementing VJ compression is called Compressed Serial Line Internet Protocol (CSLIP). See also *compression, CSLIP, SLIP*.

VLSM—Variable-Length Subnet Mask. A VLSM is a way to configure an IP subnet for maximum flexibility. Two different subnets of the same IP network number may have different masks and, therefore, different sizes. A packet is routed to the longest or most specific match. VLSM is also referred to as *Classless Inter-Domain Routing* (CIDR). See also *IP subnet, subnet mask.*

Voice-over-IP—See *VoIP*.

VoIP—Voice-over-IP. VoIP refers to set of methods for managing the transmission of voice information. The transmission takes place by means of the Internet Protocol (IP). A device can send voice data in digital form, thereby avoiding the expenses associated with ordinary telephone service. See also *IP*.

VPN—Virtual Private Network. A VPN is a private network that uses the Internet to carry all traffic. It can link all the offices, telecommuters, travelling employees, customers, and suppliers for a single organization. A VPN is virtual because it appears to the organization as a private network. Each user sees only his or her own traffic. See also *private network*.

VSA support—Vendor-Specific Attribute support. VSA support is a feature that enables companies to extend RADIUS operations without using two attributes with the same type number but different meanings.

RFC 2138, *Remote Authentication Dial In User Service (RADIUS),* specifies methods of handling vendor extensions and of encrypting and decrypting the User-Password. The RFC-defined methods differ from the way Ascend has implemented these functions in the past. In the past, Ascend extended RADIUS operations by adding Ascend vendor attributes, such as Ascend-Xmit-Rate, and used its own Ascend algorithm for User-Password encryption.

Now, the MAX ensures RADIUS RFC compliance with support for the Vendor-Specific Attribute (VSA) and the RFC-defined User-Password encryption algorithm. Ascend maintains backward compatibility by making VSA compatibility mode configurable. However, new Ascend attributes (attributes of Type 91 or smaller) are available only in VSA compatibility mode. Current Ascend attributes (attributes of Type 92 or higher) are available in both VSA compatibility mode and the default mode, which is compatible with older Ascend implementations.

RFC 2138 defines the Vendor-Specific attribute (type 26), which encapsulates attributes introduced by vendors. The older Ascend format for all attributes is as follows:

```
 0                   1                   2
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
|     Type      |    Length     | Value ...   +-+-+-+-+-+-+-+-+-+-+-
+-+-+-+-+-+-+-+-+-
```

The format of the VSA (as defined in RFC 2138) is as follows:

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|     Type      |  Length       |            Vendor-Id
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 Vendor-Id (cont)           | Vendor type   | Vendor length |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    Attribute-Specific...
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-
```

The Type of the VSA is 26. The Length is 8 or greater. Ascend's Vendor-Id is 529.

The Vendor Type, Vendor Length, and Attribute-Specific Value are the same as the Type, Length, and Value of the unencapsulated Ascend attribute found in the current dictionary. For example, the Type of the Ascend-Xmit-Rate attribute is 255. Because it is an integer, it has a Length of 6. The Value is the transmit rate of the connection. So, the fields of the VSA specify the following values:

- Type=26
- Length=12
- Vendor-Id=529
- Vendor Type=255
- Vendor Length=6
- Attribute-Specific Value=*transmit-rate*

Note: Some vendors have interpreted RFC 2138 to allow packing more than one vendor attribute in a single VSA. Ascend does not support this use. The MAX sends a single vendor attribute per VSA. If it receives a VSA that contains more than one vendor attribute, it recognizes the first vendor attribute and ignores the rest.

See also *attribute*, *RADIUS*.

VT100—An ASCII-character data terminal, consisting of a screen and keyboard. Manufactured by Digital Equipment Corporation (DEC), the VT100 has become an industry standard data terminal. VT100 emulation software allows a standard PC to act as a VT100 terminal. See also *terminal emulator*.

W

WAN—Wide Area Network. A WAN is an internet of devices, generally consisting of several networks distributed over a wide geographic distance, connected by telephone lines, and using different hardware platforms and protocol encapsulation. See also *internet*.

WAN connection—A connection between two endpoints over a WAN, as opposed to a local connection by a serial or Ethernet link. See also *WAN*.

WAN interface—The port on the MAX that is connected to a WAN line. See also *WAN*.

WAN port—A T1 or E1 port that provides a point-to-point connection between the MAX and another device. The MAX has four WAN ports. A T1 port is called a *Net/T1 port*; an E1 port is called a *Net/E1 port*.

WAN Status window—A status window that shows statistics about each active WAN link. Following is a sample WAN status window:

```
|--------------------|
|90-300 WAN Stat     |
|>Rx Pkt:     184318^|
| Tx Pkt:     159232 |
|    CRC:          0v|
|--------------------|
```

The WAN Status window shows the current count of received frames, transmitted frames, and frames with errors for each active WAN link. It also indicates the overall count for all data packets received or transmitted across the WAN. When this window is active, you can scroll down to see these three statistics for each link. The first line of each per-link count shows the name, IP address, or MAC address of the remote device.

See also *status window*.

watchdog spoofing—A method of imitating a return session-keepalive packet. An IPX server sends session keepalive packets to clients who must return the packet to keep a session active. An Ascend unit can reply to NetWare Core Protocol (NCP) watchdog packets on behalf of clients on the other side of a bridge, causing the IPX server to sense that the link is still active. Compare with *DHCP spoofing*, *IP address spoofing*, *IPX spoofing*, *SPX spoofing*.

Wide Area Network—See *WAN*.

Windows Internet Name Service—See *WINS.*

wink—On a telephone line, a signal that is comprised of an on-hook/off-hook/on-hook transition.

wink-start signaling—A signaling method in which the Customer Premises Equipment (CPE) signals an off-hook condition by sending a pulse to the Central Office (CO). Compare with *ground start signaling*, *loop start signaling.*

WINS—Windows Internet Name Service. WINS is a Microsoft product that manages the mapping between resource names and IP addresses. The Domain Name System (DNS) service used on the Internet cannot dynamically map IP addresses to local resource names. Through dynamic database updates, WINS lets a user gain access to network resources by means of user-friendly names, rather than by means of IP addresses.

wireless technology—A communications system in which electromagnetic waves carry the signal. Examples of wireless equipment include cellular telephones, pagers, the cordless mouse, and wireless transceivers for connecting to the Internet.

wiring hub—See *hub*.

X

X.3—A UTI recommendation that defines the user facilities available on all X.25 networks. See also *facility*, *X.25*.

X.3 profile—A complete set of X.3 parameters for Data Terminal Equipment (DTE) on an X.25 network. See also *DTE*, *X.25*.

X.21—A set of connector, electrical, and dialing specifications for the synchronous interface between Data Terminal Equipment (DTE) and Data Circuit-terminating Equipment (DCE) on a digital network. See also *DCE*, *DTE*.

X.21bis—An ITU-T standard that specifies the Physical-layer protocol for communication between Data Terminal Equipment (DTE) and Data Circuit-Terminating Equipment (DCE) on an X.25 network. X.21bis is nearly identical to RS-232. See also *DCE*, *DTE*, *RS-232*, *X.25*.

X.25—An international ITU-T protocol that enables users to transmit information over a packet-switched network. It allows remote devices to communicate with one another across high-speed digital links without the expense of individual nailed-up lines. The X.25 protocol handles both high-volume data transfers and interactive use of host machines. As a full-duplex, connection-oriented protocol, X.25 uses Virtual Circuits (VCs) and provides services such as multiplexing, in-sequence delivery, transfer of addressing information, segmenting and reassembly, flow control, transfer of expedited data, error control, reset, and restart. Allocation of logical channels can be either static (using a Permanent Virtual Connection, or PVC) or dynamic (using a Switched Virtual Connection, or SVC).

X.25 uses the first three layers of the OSI model. The Physical layer implements several standards, such as V.35, RS-232 and X.21bis. The Data Link layer uses an implementation of Link Access Procedure, Balanced (LAPB) and provides an error-free link between two connected devices. The Network Layer uses the Packet Layer Protocol (PLP). PLP is primarily concerned with network-routing functions and the multiplexing of simultaneous logical connections over a single physical connection.

X.25 exchanges packets between local Data Terminal Equipment (DTE) and remote Data Circuit-Terminating Equipment (DCE).

See also *DCE*, *digital modem*, *DTE*, *OSI Reference Model*, *PLP*, *PVC*, *SVC*, *X.25/PAD*, *X.25/IP*, *X.25/T3POS*.

X.25/IP—Internet Protocol over X.25. A method of transporting IP packets on X.25 facilities when the circuit is established as an end-to-end X.25 connection. See also *X.25*, *X.25/PAD*, *X.25/T3POS*.

X.25/IP inactivity timer—See *inactivity timer*.

X.25/PAD—X.25/Packet Assembler/Disassembler. In an X.25/PAD configuration, PAD-generated packets are transported using the X.25 protocol. The PAD assembles data from terminals into packets for transmission to an X.25 network, and disassembles incoming packets from the network into a separate data stream for each terminal. In addition to this multiplexing function, the PAD also provides a nearly error-free connection.

The MAX unit's X.25/PAD implementation allows users to access a public or private packet-switched network over a nailed-up ISDN connection. When a user calls X.25/PAD through a modem, the terminal server performs the authentication using a local Connection Profile or a RADIUS user profile.

See also *packet switching*, *PAD*, *X.25*, *X.25/IP*, *X.25/T3POS*.

X.25/T3POS—X.25/Transaction Processing Protocol for Point-of-Service. X.25/T3POS is a character-oriented, frame-formatted protocol designed for an X.25 packet-switched network. The protocol provides reliable and efficient data transactions between a host device and Data Terminal Equipment (DTE). The DTE is usually a client device communicating through an asynchronous port, while the host is a mainframe communicating by means of an X.25 packet network. The MAX converts data arriving from the DTE to a format capable of being transmitted over a packet network. In addition, X.25/T3POS enables you to send data over the ISDN D channel while continuing to send traffic over both B channels. See also *asynchronous transmission*, *B channel*, *D channel*, *DTE*, *X.25*, *X.25/PAD*, *X.25/IP*.

X.29—An ITU-T standard that defines the interface for the exchange of control information and user data over a packet-switched network between Data Terminal Equipment (DTE) and a Packet Assembler/Disassembler (PAD). See also *DTE*, *PAD*.

X.32—An ITU-T standard that defines the interface between Data Terminal Equipment (DTE) and Data Circuit-Terminating Equipment (DCE) for devices connecting to a public data network by means of an ISDN link, a Public Switched Telephone Network (PSTN), or a Circuit-Switched Public Data Network (CSPDN). See also *CSPDN*, *DCE*, *DTE*, *ISDN*, *PSTN*.

X.75—The ITU-T international standard for connecting packet-switched networks. See also *packet switching*.

X.121—An ITU-T standard that specifies the addressing conventions for any Data Terminal Equipment (DTE) connected to an X.25 network. See also *DTE*, *X.25*.

Xmodem—An error-correction protocol for modems. Modems that use Xmodem transmit data in 128-byte blocks. If a modem receives a block successfully, it returns a positive acknowledgment (ACK). If a modem detects an error, it sends back a negative acknowledgment (NAK) and the other modem resends the data.

Y

Yellow Alarm signal—See *RAI*.

Z

zone—An AppleTalk entity that enables you to organize the services available on your network. See also *default zone*, *zone list*.

zone list—A list of up to 32 AppleTalk zone names for the local network. Each name consists of up to 32 characters, including embedded spaces. The characters must be in the standard printing character set, and must not include an asterisk (*). See also *default zone*, *zone*.