



Configuring IP Routing

This chapter covers these topics:

- [Introduction to IP routing and interfaces](#)
- [Configuring the local IP network setup](#)
- [Configuring IP routing connections](#)
- [Configuring IP routes and preferences](#)
- [Configuring the MAX for dynamic route updates](#)
- [Managing IP routes and connections](#)

Introduction to IP routing and interfaces

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

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

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

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

Note: If you have a MAX running Multiband Simulation, IP routing is disabled.

IP addresses and subnet masks

In the MAX, you specify IP addresses in dotted decimal format (not hexadecimal). If you specify no subnet mask, the MAX assumes a default mask on the basis of address class. The default subnet mask is the default number of network bits for the address's class. [Table 10-1](#) shows the classes and

the default number of network bits for each class.

Table 10-1. IP address classes and default subnet masks

Class	Address range	Network bits
Class A	0.0.0.0 - 127.255.255.255	8
Class B	128.0.0.0 - 191.255.255.255	16
Class C	192.0.0.0 - 223.255.255.255	24

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




Figure 10-1. A class C IP address

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

IP address = 198.5.248.40

Mask = 255.255.255.248

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

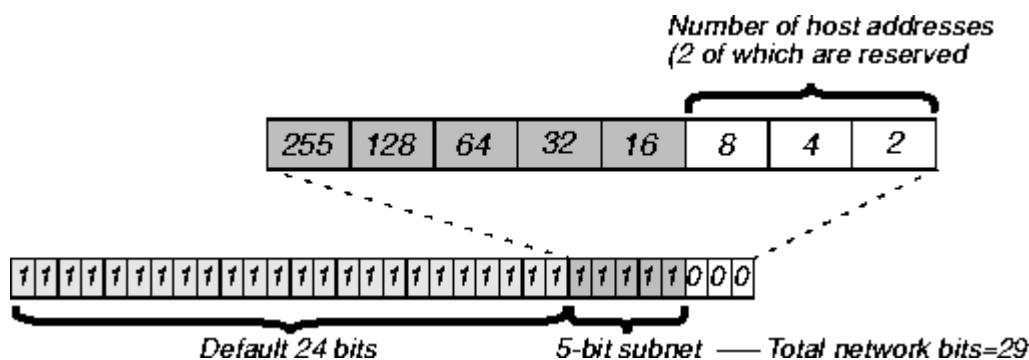


Figure 10-2. A 29-bit subnet mask and number of supported hosts

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

- 000 - Reserved for the network (base address)
- 001
- 010
- 100
- 110
- 101
- 011

111 - Reserved for the broadcast address of the subnet

Zero subnets

Early implementations of TCP/IP did not allow zero subnets. That is, subnets could have the same base address that a class A, B, or C network would have. For example, the subnet 192.168.8.0/30 was illegal because it had the same base address as the class C network 192.168.8.0/24, while 192.168.8.4/30 was legal (192.168.8.0/30 is called a zero subnet, because like a class C base address, its last octet is zero). Modern implementations of TCP/IP allow subnets to have base addresses that might be identical to the class A, B, or C base addresses. Ascend's implementations of RIP 2 and OSPF treat these so-called zero subnetworks the same as any other network. You should decide whether or not to support and configure zero subnetworks for your environment. If you configure them in some cases and treat them as unsupported in other cases, you will encounter routing problems.

[Table 10-2](#) shows how the standard subnet address format relates to Ascend notation for a class C network number.

Table 10-2. Standard subnet masks

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

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

IP address = 198.5.248.120

Mask = 255.255.255.248

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

198.5.248.120 - 198.5.248.127

A host route is a special case IP address with a subnet mask of 32 bits. It has a subnet mask of 255.255.255.255.

IP routes

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

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

How the MAX uses the routing table

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

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

Static and dynamic routes

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

- ! Each Static Rtes profile specifies one static route. If a path to a destination must be reliable, the administrator often configures more than one path (a secondary route), in which case the MAX chooses the route based on assigned metrics and availability.
- ! The Ethernet > Mod Config profile specifies a static connected route, which states "to reach system-A, send packets out this interface to system-A." Connected routes are low cost, because no remote connection is involved.
- ! Each IP-routing Connection profile specifies a static route that states *to reach system-A, send packets out this interface to system-B*, where system-B is another router.

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

Route preferences and metrics

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

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

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

Note: You can configure the DownMetric and DownPreference parameters to assign different metrics or preferences to routes on the basis of whether the route is in use or is down. You might want to direct the MAX to use active routes, if available, rather than choose routes that are down.

MAX IP interfaces

Ethernet interfaces

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

```
** Ascend MAX Terminal Server **
```

```
ascend% iproute show
```

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

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

Following are descriptions of the interfaces created at startup:

- ! The Ethernet IP interface, labeled `ie0`, is always active, because it is always connected. Its IP address is assigned in Ethernet > Mod Config > Ether Options.
- The MAX creates two routing table entries: one with a destination of the network (labeled `ie0`), and the other with a destination of the MAX (labeled `local`).
- ! The black-hole (`bh0`) interface is always up. The black-hole address is 127.0.0.3. Packets routed to this interface are discarded silently.
 - ! The loopback (labeled `local`) interface is always up. The loopback address is 127.0.0.1/32.
 - ! The reject (labeled `rj0`) interface is always up. The reject address is 127.0.0.2. Packets routed to this interface are sent back to the source address with an ICMP "host unreachable" message.
 - ! Multicast interfaces have a destination address with a value of 224 for the first octet. For information about multicast addresses, see [Chapter 12, Setting Up IP Multicast Forwarding](#).
 - ! Not shown in the example is an inactive interface. It is created when you configure a Connection profile. The inactive interface is where all routes point when their WAN connections are down. The inactive interface label is `wanidle0`.

WAN IP interfaces

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

Numbered interfaces

The MAX can operate as both a system-based and interface-based router. Interface-based routing uses numbered interfaces. Some routers or applications require numbered interfaces, and some sites use them for trouble-shooting leased point-to-point connections and forcing routing decisions between two links going to the same final destination. More generally, interface-based routing allows the MAX to operate in much the same way as a multihomed Internet host.

[Figure 10-3](#) shows a sample interface-based routing connection.




Figure 10-3. Interface-based routing example

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

With system-based routing, these addresses are not assigned. The site A MAX routes packets to the

remote network on the basis of the WAN interface it created when the connection was brought up, rather than a configured IP address.

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

- ! Connections > IP Options > IF Adrs (the link-specific address for the MAX)
- ! Connections > IP Options > WAN Alias (the far end link-specific address)

Or, you may omit the remote side's system-based IP address from the Connection profile and use interface-based routing exclusively. This is an appropriate mechanism, for example, if the remote system is on a backbone net that might be periodically reconfigured by its administrators, and you want to refer to the remote system only by its mutually agreed-upon interface address. In this case, the link-specific IP addresses are specified in the following parameters:

- ! Connections > IP Options > IF Adrs (the near end numbered interface)
- ! Connections > IP Options > LAN Adrs (the far end numbered interface)

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

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

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

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

Configuring the local IP network setup

The Ethernet profile configures system-global parameters that affect all IP interfaces in the MAX. These are the related parameters:

```
Ethernet
  Mod Config
    Ether options
      IP Adrs=10.2.3.1/24
      2nd Adrs=0.0.0.0/0
      RIP=Off
      Ignore Def Rt=Yes
```

Proxy Mode=Off

WAN options...

Pool#1 start=100.1.2.3
Pool#1 count=128
Pool#1 name=Engineering Dept.
Pool#2 start=0.0.0.0
Pool#2 count=0
Pool#2 name=
Pool#3 start=10.2.3.4
Pool#3 count=254
Pool#3 name=Marketing Dept.
Pool#4 start=0.0.0.0
Pool#4 count=0
Pool#4 name=
Pool#5 start=0.0.0.0
Pool#5 count=0
Pool#5 name=
Pool#6 start=0.0.0.0
Pool#6 count=0
Pool#6 name=
Pool#7 start=0.0.0.0
Pool#7 count=0
Pool#7 name=
Pool#8 start=0.0.0.0
Pool#8 count=0
Pool#8 name=
Pool#9 start=0.0.0.0
Pool#9 count=0
Pool#9 name=
Pool#A start=0.0.0.0
Pool#A count=0
Pool#A name=
Pool only=No
Pool Summary=No

Shared Prof=No

Telnet PW=Ascend

BOOTP Relay...

BOOTP Relay Enable=No
Server=N/A
Server=N/A

DNS...

Domain Name=abc.com
Sec Domain Name=
Pri DNS=10.65.212.10
Sec DNS=12.20 7.23.51
Allow As Client DNS=Yes
Pri WINS=0.0.0.0
Sec WINS=0.0.0.0
List Attempt=No
List Size=N/A
Client Pri DNS=0.0.0.0
Client Sec DNS=0.0.0.0

SNTP Server...

SNTP Enabled=Yes
Time zone-UTC+0000
SNTP host#1=0.0.0.0
SNTP host#2=0.0.0.0
SNTP host#3=0.0.0.0

```
UDP Cksum=No
Adv Dialout Routes=Always
```

Understanding the IP network parameters

This section provides some background information about the IP network configuration. The information is organized by functionality rather than by parameter. For more information on each parameter, see the *MAX Reference Guide*.

Primary IP address for each Ethernet interface

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

Second IP address for each Ethernet interface

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

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

Dual IP is also used to distribute the load of routing traffic to a large subnet, by assigning IP addresses on that subnet to two or more routers on the backbone. When the routers have a direct connection to the subnet as well as to the backbone network, they route packets to the subnet and include the route in their routing table updates.

Dual IP also allows you to make a smooth transition when changing IP addresses. That is, a second IP address can act as a placeholder while you are making the transition in other network equipment.

[Figure 10-4](#) shows an example IP network to which a MAX is connected:




Figure 10-4. Sample dual IP network

Two IP addresses are assigned to each of the MAX's Ethernet interfaces. 10.1.2.4 and 11.6.7.9 are assigned to Ethernet 1. 12.1.1.2 and 13.9.7.5 are assigned to Ethernet 2. In this example, the MAX routes between all displayed networks. The MAX enables the host assigned 12.1.1.1 to communicate with the host assigned 13.9.7.4 and the host assigned 10.1.2.3.

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

Enabling RIP on the Ethernet interface

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

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

Ignoring the default route

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

Proxy ARP and inverse ARP

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

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

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

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

Specifying address pools

You can define up to 10 address pools in the Ethernet profile, with each pool supporting up to 254 addresses. The Pool#N start parameter specifies the first address in a block of contiguous addresses on the local network or subnet. The Pool#N count parameter specifies how many addresses are in the pool (up to 255). Addresses in a pool do not accept a netmask modifier, because they are advertised as host routes. If you allocate IP addresses on a separate IP network or subnet, make sure you inform other IP routers about the route to that network or subnet, either by statically configuring those routes or configuring the MAX to dynamically send updates.

Forcing callers configured for a pool address to accept dynamic assignment

During PPP negotiation, a caller may reject the IP address offered by the MAX and present its own IP address for consideration. Connection profiles compare IP addresses as part of authentication, so the MAX would automatically reject such a request if the caller has a Connection profile. However, Name-password profiles have no such authentication mechanism, and could potentially allow a caller to spoof a local address. The Pool Only parameter instructs the MAX to hang up if a caller rejects the dynamic assignment.

Summarizing host routes in routing table advertisements

IP addresses assigned dynamically from a pool are added to the routing table as individual host routes. You can summarize this network (the entire pool), cutting down significantly on route flapping and the size of routing table advertisements.

Pool Summary indicates the route summarization is in use; that is, a series of host routes will be summarized into a network route advertisement. Packets destined for a valid host address on that network are routed to the host, and packets destined for an invalid host address are rejected with an ICMP *host unreachable* message.

To use the pool summary feature, create a network-aligned pool and set the Pool Summary parameter to Yes. To be network-aligned, the Pool Start address must be the first host address. Subtract one from the Pool Start address to determine the network address (the zero address on the subnet). Since the first and last address of a subnet are reserved, you must set the Pool Count to a value that is 2 less than a power of 2. For example, you may use values 2, 6, 14, 30, 62, 126 or 253. The netmask will be deduced from a value that is 2 greater than Pool Count. For example, with this configuration:

```
Pool Summary=Yes
Pool#1 start=10.12.253.1
Pool#1 count=126
```

The network alignment address is Pool Start address -1: 10.12.253.0 and the netmask is Pool Count +2 addresses: 255.255.255.128. The resulting address pool network is:

```
10.12.253.0/25
```

For an example configuration that shows route summarization, see [Configuring DNS](#).

! Sharing Connection profiles

The Shared Prof parameter specifies whether the MAX will allow more than one incoming call to share the same Connection profile. This feature is related to IP routing because sharing profiles cannot result in two IP addresses reached through the same profile.

In low-security situations, more than one dial-in user can share a name and password for accessing the local network. This would require sharing a single Connection profile that specifies bridging only, or dynamic IP address assignment. Each call would be a separate connection. The name and password would be shared, and a separate IP address would be assigned dynamically to each caller.

If a shared profile uses an IP address, it must be assigned dynamically, because multiple hosts cannot share a single IP address.

Telnet password

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

BOOTP Relay

By default, a MAX does not relay BOOTP (Bootstrap Protocol) requests to other networks. If BOOTP is enabled, the MAX can relay BOOTP requests to another network. However, SLIP BOOTP must be disabled in Ethernet > Mod Config > TServ Options. SLIP BOOTP makes it possible for a computer connecting to the MAX over a SLIP connection to use the Bootstrap Protocol. A MAX can support BOOTP on only one connection. If both SLIP BOOTP and BOOTP relay are enabled, you will receive an error message.

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

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

Local domain name

The Domain Name is used for DNS lookups. When the MAX is given a hostname to look up, it tries various combinations including appending the configured domain name. The secondary domain name (Sec Domain Name) can specify another domain name that the MAX can search using DNS. The MAX searches the secondary domain only after the domain specified in the Domain Name parameter.

DNS or WINS name servers

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

DNS lists

DNS can return multiple addresses for a hostname in response to a DNS query, but it does not include information about availability of those hosts. Users typically attempt to access the first address in the list. If that host is unavailable, the user must try the next host, and so forth. However, if the access attempt occurs automatically as part of immediate services, the physical connection is torn down when the initial connection fails. To avoid tearing down physical links when a host is unavailable, you can use the List Attempt parameter to enable the user to try one entry in the DNS list of hosts, and if that connection fails, to try the next entry, and so on, without losing the WAN session. The List Size parameter specifies the maximum number of hosts listed (up to 35).

Client DNS

Client DNS configurations define DNS server addresses that will be presented to WAN connections during IPCP negotiation. They provide a way to protect your local DNS information from WAN users. Client DNS has two levels: a global configuration that applies to all PPP connections (defined in the Ethernet profile), and a connection-specific configuration that applies only to the WAN connection defined in the Connection profile. The global client addresses are used only if none are specified in the Connection profile.

SNTP service

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

```
UTC +0130
```

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

```
UTC +0800
```

For Frankfurt, which is 1 hour behind UTC:

```
UTC -0100
```

Specifying SNTP server addresses

The host parameter lets you specify up to three server addresses. The MAX attempts to communicate with the first address. It attempts the second only if the first is inaccessible, and the third only if the second is inaccessible.

UDP checksums

If data integrity is of the highest concern for your network and having redundant checks is important, you can turn on UDP checksums to generate a checksum whenever a UDP packet is transmitted. UDP packets are transmitted for queries and responses related to ATMP, SYSLOG, DNS, ECHOSERV, RADIUS, TACACS, RIP, SNTP, and TFTP.

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

Poisoning dialout routes in a redundant configuration

If you have another Ascend unit backing up the MAX in a redundant configuration on the same network, you can use the Adv Dialout Routes parameter to instruct the MAX to stop advertising IP routes that use dial services if its trunks are in the alarm condition. Otherwise, it continues to advertise its dialout routes, which prevents the redundant unit from taking over the routing responsibility.

Example IP network configurations

This section shows some example Ethernet profile IP configurations. For a more complete example that shows an Ethernet profile, Route profile, and Connection profile configuration that work together, see [Configuring DNS](#).

Configuring the MAX IP interface on a subnet

On a large corporate backbone, many sites configure subnets to increase the network address space, segment a complex network, and control routing in the local environment. For example, [Figure 10-5](#) shows the main backbone IP network (10.0.0.0) supporting an Ascend GRF router (10.0.0.17):




Figure 10-5. Creating a subnet for the MAX

You can place the MAX on a subnet of that network by entering a subnet mask in its IP address specification, for example:

1. Open Ethernet > Mod Config > Ether Options.
2. Specify the IP subnet address for the MAX on Ethernet. For example:

```
Ethernet
  Mod Config
    Ether options
      IP Adrs=10.2.3.1/24
```

3. Configure the MAX to receive RIP updates from the local GRF router.

```
RIP=Recv=v2
```

4. Close the Ethernet profile.

With this subnet address, the MAX requires a static route to the backbone router on the main network; otherwise, it can only communicate with devices on the subnets to which it is directly connected. To create the static route and make the backbone router the default route:

1. Open the Default IP Route profile.
2. Specify the IP address of a backbone router in the Gateway parameter. For example:

```
Ethernet
  Static Rtes
    Name=Default
    Active=Yes
    Dest=0.0.0.0/0
    Gateway=10.0.0.17
    Preference=100
    Metric=1
    DownPreference=140
    DownMetric=7
    Private=Yes
```

3. Close the Default IP Route profile.

See [Configuring IP routes and preferences](#) for more information about IP Route profiles. To verify that the MAX is up on the local network, invoke the terminal server interface and enter the Ping command to a local IP address or hostname. For example:

```
ascend% ping 10.1.2.3
```

You can terminate the Ping exchange at any time by typing Ctrl-C.

Configuring DNS

The DNS configuration enables the MAX to use local DNS or WINS servers for lookups. In this example DNS configuration, client DNS is not in use. Note that you can protect your DNS servers from callers by defining connection-specific ("client") DNS servers and specifying that Connection profiles use those client servers. To configure the local DNS service:

1. Open Ethernet > Mod Config > DNS.
2. Specify the local domain name.
3. If appropriate, specify a secondary domain name.
4. Specify the IP addresses of a primary and secondary DNS server, and turn on the DNS list attempt feature.

```
Ethernet
  Mod Config
    DNS...
      Domain Name=abc.com
      Sec Domain Name=
      Pri DNS=10.65.212.10
      Sec DNS=12.20 7.23.51
      Allow As Client DNS=Yes
      Pri WINS=0.0.0.0
      Sec WINS=0.0.0.0
      List Attempt=Yes
      List Size=35
      Client Pri DNS=0.0.0.0
      Client Sec DNS=0.0.0.0
      Enable Local DNS Table=No
      Loc.DNSTab Auto Update=No
```

5. Close the Ethernet profile.

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

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

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

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

- ! # Reads (the number of reads since entry was created)

This field is updated each time a local name query match is found in the local DNS table.

- ! Time of Last Read

You can check the list of host names and IP addresses in the table using the `termserv` command `show dnstab`. [Figure 10-6](#) shows an example of a DNS table on a MAX. Other terminal server commands show individual entries, with a list of IP addresses for the entry.

Local DNS Table

Name	IP Address	# Reads	Time of last read
1: ""	-----	-----	
2: "server.corp.com."	200.0.0.0	2	Feb 10 10:40:44
3: "boomerang"	221.0.0.0	2	Feb 10 9:13:33
4: ""	-----	-----	
5: ""	-----	-----	
6: ""	-----	-----	
7: ""	-----	-----	

Figure 10-6. Local DNS table example

New terminal server command changes

New `show` and `dnstab` commands have been added to help you view, edit, or make entries in the DNS table.

show commands

- ! `show ?` displays a list that includes `dnstab` help.
- ! **show dnstab** displays the local DNS table.
- ! **show dnstab ?** displays help for the `dnstab` editor.
- ! **show dnstab entry** displays the local DNS table entry (all IP addresses in the list)

dnstab commands

The terminal server `dnstab` command has these variations:

Table 10-3. dnstab commands

dnstab Command	Description
<code>dnstab</code>	Displays help information about the DNS table.
<code>dnstab show</code>	Displays the local DNS table.
	Displays a list for entry <i>n</i> in the local DNS table.

<code>dnstab entry n</code>	The list displayed includes the entry and all the IP addresses stored for that entry up to a maximum number of entries specified in the List Size parameter. If List Attempt=No, no list is displayed.
<code>dnstab edit</code>	Start editor for the local DNS table.

Configuring the local DNS table

To enable and configure the local DNS table:

1. Display Ethernet Profile: Ethernet > Mod Config > DNS menu.
2. Select a setting for the List Attempt parameter.
3. Specify the list size by setting the List Size parameter.
4. Select Enable Local DNS Table=Yes.

The default is No.

5. Select a setting for the Loc.DNS Tab Auto Update parameter.

Criteria for valid names in the local DNS table

- ! Must be unique in the table.
- ! Must start with an alphabetic character, which may be either upper- or lower-case.
- ! Must be less than 256 characters
- ! Names can be local names or fully qualified names that include the domain name.

Periods at the end of names are ignored.

Entering IP addresses in the local DNS table

To enter IP addresses in a local DNS table, you use the DNS table editor from the terminal server. While the editor is in use, the system cannot look up addresses in the table or perform automatic updates. A table *entry* is one of the eight table indexes. It includes the host name, IP address (or addresses), and information fields. To place the initial entries in the table:

1. At the terminal server interface, type `dnstab edit`.

Before you make any entries, the table is empty. The editor initially displays zeros for each of the eight entries in the table. To exit the table editor without making an entry, press Enter.

2. Type an entry number and press Enter.

A warning appears if you type an invalid entry number. If the entry exists, the current name for that entry appears in the prompt.

3. Type the name for the current entry.

If the system accepts the name, it places the name in the table and prompts you for the IP address for the name that you just entered. (For the characteristics of a valid name, see [Criteria for valid names in the local DNS table](#).)

If you enter an invalid name, the system prompts you to enter a valid name.

4. Type the IP address for the entry.

If you enter an address in the wrong format, the system prompts you for the correct format. If your format is correct, the system places the address in the table and the editor prompts you for the next entry.

5. When you are finished making entries, type the letter O and press Enter when the editor prompts you for another entry.

Editing the local DNS table

To edit the DNS table entries, you access the DNS table editor from the terminal server. While the editor is in use, the system cannot look up addresses in the table or perform automatic updates. A table *entry* is one of the eight table indexes. It includes the host name, IP address (or addresses), and information fields. To edit one or more entries in the local DNS table:

1. At the terminal server interface, type `dnstab edit`.

If the table has already been created, the number of the entry last edited appears in the prompt.

2. Type an entry number or press Enter to edit the entry number currently displayed.

A warning appears if you type an invalid entry number. If the entry exists, the current value for that entry appears in the prompt.

3. Replace, accept, or clear the displayed name, as follows:

- To replace the name, type a new name and press Enter.

- To accept the current name, press Enter.

- To clear the name, press the spacebar and then Enter.

If you enter a valid name, the system places it in the table (or leaves it there if you accepted the current name) and prompts you for the corresponding IP address. (For the characteristics of a valid name, see [Criteria for valid names in the local DNS table](#).)

If you clear an entry name, all information in all fields for that entry is discarded.

4. Either type a new IP address and press Enter, or leave the current address and just press Enter.

- If you are changing the name of the entry but not the IP address, press Return.

- To change the IP address, type the new IP address

If the address is in the correct format, the system places it in the table and prompts you for another entry.

- When you are finished making entries, type the letter O and press Enter when the editor prompts you for another entry.

Deleting an entry from the local DNS table

To delete an entry from the local DNS table:

- At the terminal server interface, type `dnstab edit` to display the table.
- Type the number of the entry you want to delete and press Return.
- Press the spacebar and then press Return.

Setting up address pools with route summarization

The address pool parameters enable the MAX to assign an IP address to incoming calls that are configured for dynamic assignment. These addresses are assigned on a first-come first-served basis. After a connection has been terminated, its address is freed up and returned to the pool for reassignment to another connection. [Figure 10-7](#) shows a host using PPP dial-in software to connect to the MAX.




Figure 10-7. Address assigned dynamically from a pool

This example shows how to set up network-aligned address pools and use route summarization. It also shows how to enter a static route for the pool subnet and make Connection profile route private, which are requirements when using route summarization.

These are the rules for network-aligned address pools:

- The Pool Count must be two less than the total number of addresses in the pool.
Add two to Pool Count for the total number of addresses in the subnet, and calculate the netmask for the subnet based on this total.
- The Pool Start address must be the first host address.
Subtract 1 from the Pool Start address for the base address for the subnet.

For example, the following configuration is network aligned:

```
Ethernet
  Mod Config
    WAN options...
      Pool#1 start=10.12.253.1
```

```
Pool#1 count=62
Pool#1 name=Engineering Dept.
Pool Summary=Yes
```

Pool Start is set to 10.12.253.1. When you subtract one from this address, you get 10.12.253.0, which is a valid base address for the 255.255.255.192 netmask. Note that 10.12.253.64, 10.12.253.128, and 10.12.253.192 are also valid zero addresses for the same netmask. The resulting address pool network is 10.12.253.0/26.

Pool Count is set to 62. When you add two to the Pool Count, you get 64. The netmask for 64 addresses is 255.255.255.192 (256-64 = 192). The Ascend subnet notation for a 255.255.255.192 netmask is /26.

After verifying that *every one* of the configured address pools is network-aligned, you must enter a static route for them. These static routes handle all IP address that have not been given to users by routing them to the reject interface or the blackhole interface. (See [MAX IP interfaces.](#))

Note: The MAX creates a host route for every assigned address from the pools and host routes override subnet routes. So, packets whose destination matches an assigned IP address from the pool are properly routed and not discarded or bounced. Because the MAX advertises the entire pool as a route, and only privately knows which IP addresses in the pool are active, a remote network might improperly send the MAX a packet to an inactive IP address. Depending on the static route specification, these packets are either bounced with an ICMP unreachable or silently discarded.

For example, the following static route specifies the blackhole interface, so it silently discards all packets whose destination falls in the pool's subnet. In addition to the Dest and Gateway parameters that define the pool, be sure you have set the Metric, Preference, Cost, and Private parameters as shown.

```
Ethernet
  Static Rtes
    Name=pool-net
    Active=Yes
    Dest=10.12.253.0/26
    Gateway=127.0.0.0
    Preference=0
    Metric=0
    Cost=0
    Private=No
```

The routing table will contain the following lines:

Destination	Gateway	IF	Flg	Pref	Met	Use	Age
10.12.253.0/26	-	bh0	C	0	0	0	172162
127.0.0.0/32	-	bh0	CP	0	0	0	172163
127.0.0.1/32	-	lo0	CP	0	0	0	172163
127.0.0.2/32	-	rj0	CP	0	0	0	172163

When you configure Connection profiles that assign IP addresses from the pool, make sure the Private parameter is set to Yes. For example:

```

Ethernet
  Connections
    Ip options...
      LAN Adrs=0.0.0.0/0
      WAN Alias=0.0.0.0
      IF Adrs=0.0.0.0/0
      Preference=100
      Cost=0
      Private=Yes
      RIP=Off
      Pool=1

```

Configuring IP routing connections

When IP routing is enabled and addresses are specified in a Connection profile, it defines an IP WAN interface. These are the related options:

```

Ethernet
  Answer
    Assign Adrs=Yes
    PPP options...
      Route IP=Yes

    Session options...
      RIP=Off

Ethernet
  Connections
    Station=remote-device

    Route IP=Yes
    IP options...
      LAN Adrs=0.0.0.0/0
      WAN Alias=0.0.0.0/0
      IF Adrs=0.0.0.0/0
      Preference=100
      Metric=7
      DownPreference=120
      DownMetric=9
      Private=No
      RIP=Off
      Pool=0

    Session options...
      IP Direct=0.0.0.0

```

Understanding the IP routing connection parameters

This section provides some background information about enabling IP routing in the Answer profile and Connection profiles. For more information on each parameter, see the *MAX Reference Guide*.

Enabling dynamic address assignment for answered calls

Assign Adrs must be set to Yes in the Answer profile to enable the MAX to allocate IP addresses dynamically from a pool of designated addresses on the local network. The caller's PPP software must be configured to accept an address dynamically. If the Pools Only parameter is set to Yes in the Ethernet profile, the MAX terminates connections that reject the assigned address during PPP negotiation. See [Configuring dynamic address assignment to a dial-in host](#) for related information.

Enabling IP routing for WAN connections

Route IP in Answer > PPP Options must be set to Yes to enable the MAX to negotiate a routing connection.

Enabling IP routing for a WAN interface

To enable IP packets to be routed for this connection, set the Route IP parameter to Yes in the Connection profile. When IP routing is enabled, IP packets are always routed, they are never bridged.

Configuring the remote IP address

The LAN parameter specifies the IP address of the remote device. Before accepting a call from the far end, the MAX matches this address to the source IP address presented by the calling device. It may be one of the following values:

- ! IP address of a router

If the remote device is an IP router, specify its address including its netmask modifier. (See [IP addresses and subnet masks](#) for background information.) If you omit the netmask, the MAX inserts a default netmask which makes the entire far-end network accessible.

- ! IP address of a dial-in host

If the remote device is a dial-in host running PPP software, specify its address including a netmask modifier of /32; for example, 10.2.3.4/32.

- ! The null address (0.0.0.0)

If the remote device is a dial-in host that will accept dynamic address assignment, leave the remote-address parameter blank.

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

WAN alias

This is another IP address for the remote device, used for numbered interface routing. The WAN Alias will be listed in the routing table as a gateway (next hop) to the Lan Adrs. The caller must be using a numbered interface, and its interface address must agree with the WAN Alias setting.

Specifying a local IP interface address

This is another local IP interface address, to be used as the local numbered interface instead of the default (the Ethernet IP Adrs).

Assigning metrics and preferences

Connection profiles often represent switched connections, which have an initial cost that can be avoided if a nailed-up link to the same destination can be used. To favor nailed-up links, you can assign a higher metric to switched connections than any of the nailed-up links that can go to the same place.

Each connection represents a static route, which has a default preference of 100. (See [Route preferences and metrics](#).) For each connection, you can fine-tune the route preference and assign a different preference.

Note: You can configure the DownMetric and DownPreference parameters to assign different metrics or preferences to routes on the basis of whether the route is in use or is down. You might want to direct the MAX to use active routes, if available, rather than choose routes that are down.

Private routes

The Private parameter specifies whether the MAX discloses the existence of this route when queried by RIP or another routing protocol. The MAX uses Private routes internally; they are not advertised.

Assigning the IP address dynamically

The Pool parameter specifies an IP address pool from which the caller will be assigned an IP address. If the Pool parameter is null but all other configuration settings enable dynamic assignment, the MAX gets IP addresses from the first defined address pool. See [Configuring DNS](#).

IP direct configuration

An IP Direct configuration bypasses routing and bridging tables for all incoming packets and sends each packet received to the specified IP address. All outgoing packets are treated as normal IP traffic. They are not affected by the IP Direct configuration.

Note: IP Direct connections are typically configured with RIP turned off. If you set the IP Direct configuration with RIP set to receive, all RIP updates will be forwarded to the specified address. This is typically not desirable since RIP updates are designed to be stored locally by the IP router (the MAX, in this case).

Configuring RIP on this interface

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

Ascend recommends that you run RIP version 2 (RIP-v2) if possible. Ascend does not recommend running RIP-v2 and RIP-v1 on the same network in such a way that the routers receive each other's advertisements. RIP-v1 does not propagate subnet mask information, and the default class network mask is assumed, while RIP-v2 handles subnet masks explicitly. Running the two versions on the same network can result in RIP-v1 *guesses* overriding accurate subnet information obtained via RIP-v2.

Checking remote host requirements

IP hosts, such as UNIX systems, Windows or OS/2 PCs, or Macintosh systems, must have appropriately configured TCP/IP software. A remote host calling into the local IP network must also have PPP software.

UNIX

UNIX systems typically include a TCP/IP stack, DNS software, and other software, files, and utilities used for Internet communication. UNIX network administration documentation describes how to configure these programs and files.

Window or OS/2 software

PCs running Windows or OS/2 need the TCP/IP networking software. The software is included with Windows 95, but the user may need to purchase and install it separately if the computer has a previous version of Windows or OS/2.

Macintosh software

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

Software configuration

For any platform, the TCP/IP software must be configured with the host's IP address and subnet mask. If the host will obtain its IP address dynamically from the MAX, the TCP/IP software must be configured to allow dynamic allocation. If a DNS server is supported on your local network, you should also configure the host software with the DNS server's address.

Typically, the host software is configured with the MAX as its default router.

Examples IP routing connections

This section provides example Connection profile configurations for IP routing. These examples all presume that the Ethernet profile has been configured correctly, as described in [Configuring the local IP network setup](#).

Configuring dynamic address assignment to a dial-in host

In this example, the dial-in host is a PC that will accept an IP address assignment from the MAX dynamically. [Figure 10-8](#) shows an example network.




Figure 10-8. A dial-in user requiring dynamic IP address assignment

In this example, site A is a backbone network and site B is a single dial-in host with a modem, TCP/IP stack, and PPP software. The PPP software running on the PC at site B must be configured to acquire its IP address dynamically. For example, this example software configuration presumes that the PC has a modem connection to the MAX:

```

Username=victor
Accept Assigned IP=Yes
IP address=Dynamic (or Assigned or N/A)
Netmask=255.255.255.255 (or None or N/A)
Default Gateway=None or N/A
Name Server=10.2.3.55
Domain suffix=abc.com
Baud rate=38400
Hardware handshaking ON
VAN Jacobsen compression ON

```

To configure the MAX to accept dial-in connections from site B and assign an IP address:

1. Open Ethernet > Mod Config > WAN Options.
2. Type the start address of the pool and the number of contiguous addresses it includes. For example:

```

Ethernet
  Mod Config
    WAN options
      Pool#1 start=10.12.253.1
      Pool#1 count=126
      Pool#1 name=Engineering Dept.
      Pool only=Yes
      Pool Summary=Yes

```

3. Open the Ether Options subprofile and turn on Proxy Mode.

```

Ether options
  Proxy Mode=Yes

```

4. Close the Ethernet profile.
5. Open the Answer profile and enable both IP routing and dynamic address assignment.

```

Ethernet
  Answer
    Assign Adrs=Yes
    PPP options
      Route IP=Yes

```

6. Close the Answer profile.
7. Open a Connection profile for the dial-in user.
8. Specify the user's name, activate the profile, and set encapsulation options.

```

Ethernet
  Connections
    Station=victor
    Active=Yes
    Encaps=PPP
    Encaps options...
      Send Auth=CHAP
      Recv PW=*SECURE*

```

9. Configure IP routing and address assignment.

```

Route IP=Yes
IP options
  LAN Adrs=0.0.0.0/0
  RIP=Off
  Pool=1

```

10. Close the Connection profile.

Configuring a host connection with a static address

This type of connection enables the dial-in host to keep its own IP address when logging into the MAX IP network. For example, if a PC user telecommutes to one IP network and uses an ISP on another IP network, one of those connections can assign an IP address dynamically and the other can configure a host route to the PC. This example shows how to configure a host connection with a static address. See [IP addresses and subnet masks](#) for details on the /32 netmask.




Figure 10-9. A dial-in user requiring a static IP address (a host route)

In this example, the PC at site B is running PPP software that includes settings like these:

```

Username=patti
Accept Assigned IP=N/A (or No)
IP address=10.8.9.10
Netmask=255.255.255.255
Default Gateway=N/A (or None)
Name Server=10.7.7.1
Domain suffix=abc.com
VAN Jacobsen compression ON

```

To configure the MAX to accept dial-in connections from site B:

1. Open the Answer profile and enable IP routing.

```

Ethernet
  Answer
    PPP options
      Route IP=Yes

```

2. Close the Answer profile.
3. Open a Connection profile for the dial-in user.
4. Specify the user's name, activate the profile, and set encapsulation options.

```

Ethernet
  Connections
    Station=patti
    Active=Yes

```

```

Encaps=PPP
Encaps options...
  Send Auth=CHAP
  Recv PW=*SECURE*

```

5. Configure IP routing.

```

Route IP=Yes
IP options
  LAN Adrs=10.8.9.10/32
  RIP=Off

```

6. Close the Connection profile.

Configuring an IP Direct connection

You can configure a Connection profile to automatically redirect incoming IP packets to a specified host on the local IP network without having the packets pass through the routing engine on the MAX.




Figure 10-10. Directing incoming IP packets to one local host

Note: IP Direct connections typically turn off RIP. If the connection is configured to receive RIP, all RIP packets from the far side are kept locally and forwarded to the IP address you specify for IP Direct.

To configure an IP Direct connection:

1. Open the Answer profile and enable IP routing.

```

Ethernet
  Answer
    PPP options
      Route IP=Yes

```

2. Close the Answer profile.

3. Open a Connection profile for the dial-in connection.

4. Specify the remote device's name, activate the profile, and set encapsulation options.

```

Ethernet

```

```

Connections
  Station=Pipeline1
  Active=Yes
  Encaps=MPP
  Encaps options...
    Send Auth=CHAP
    Recv PW=localpw
    Send PW=remotepw

```

5. Configure IP routing.

```

Route IP=Yes
IP options
  LAN Adrs=10.8.9.10/22
  RIP=Off

```

6. Open the Session Options subprofile and specify the IP Direct host.

```

Session options
  IP Direct=10.2.3.11

```

7. Close the Connection profile.

Note: The IP Direct address you specify in Connections > Session Options is the address to which all incoming packets on this connection will be directed. When you use the IP Direct feature, a user cannot Telnet directly to the MAX from the far side. All incoming IP traffic is directed to the specified address on the local IP network.

Configuring a router-to-router connection

In this example, the MAX is connected to a corporate IP network and needs a switched connection to another company that has its own IP configuration. [Figure 10-11](#) shows an example network diagram.




Figure 10-11. A router-to-router IP connection

This example assumes that the Answer profile in both devices enable IP routing. To configure the site A MAX for a connection to site B:

1. Open a Connection profile for the site B device.
2. Specify the remote device's name, activate the profile, and set encapsulation options.

```

Ethernet
  Connections

```

```
Station=PipelineB
Active=Yes
Encaps=MPP
Encaps options...
  Send Auth=CHAP
  Recv PW=localpw
  Send PW=remotepw
```

3. Configure IP routing.

```
Route IP=Yes
IP options
  LAN Adrs=10.9.8.10/22
  RIP=Off
```

4. Close the Connection profile.

To configure the site B Pipeline:

1. Open the Connection profile for the site A MAX.
2. Specify the site A MAX unit's name, activate the profile, and set encapsulation options.

```
Ethernet
  Connections
    Station=MAXA
    Active=Yes
    Encaps=MPP
    Encaps options...
      Send Auth=CHAP
      Recv PW=localpw
      Send PW=remotepw
```

3. Configure IP routing.

```
Route IP=Yes
IP options
  LAN Adrs=10.2.3.1/22
  RIP=Off
```

4. Close the Connection profile.

Configuring a router-to-router connection on a subnet

In this example network, the MAX is used to connect telecommuters with their own Ethernet networks to the corporate backbone. The MAX is on a subnet, and assigns subnet addresses to the telecommuters' networks.




Figure 10-12. A connection between local and remote subnets

This example assumes that the Answer profile in both devices enables IP routing. Because the MAX specifies a netmask as part of its own IP address, the MAX must use other routers to reach IP addresses outside that subnet. To forward packets to other parts of the corporate network, the MAX must either have a default route configuration to a router in its own subnet (such as the Cisco router in Figure 5-12) or it must enable RIP on Ethernet.

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

1. Open a Connection profile for the site B device.
2. Specify the remote device's name, activate the profile, and set encapsulation options.

```
Ethernet
  Connections
    Station=PipelineB
    Active=Yes
    Encaps=MPP
    Encaps options...
      Send Auth=CHAP
      Recv PW=localpw
      Send PW=remotepw
```

3. Configure IP routing.

```
Route IP=Yes
IP options
  LAN Adrs=10.7.8.200/24
  RIP=Off
```

4. Close the Connection profile.

To specify the local Cisco router as the MAX unit's default route:

1. Open the Default IP Route profile.
2. Specify the Cisco router's address as the gateway address.

```
Ethernet
  Static Rtes
    Name=Default
    Active=Yes
    Dest=0.0.0/0
```

```

Gateway=10.4.4.133
Metric=1
Preference=10
Private=Yes

```

3. Close the IP Route profile.

To configure the site B Pipeline unit for a connection to site A:

1. Open the Connection profile in the Pipeline unit for the site A MAX.
2. Specify the site A MAX unit's name, activate the profile, and set encapsulation options.

```

Ethernet
  Connections
    Station=MAXA
    Active=Yes
    Encaps=MPP
    Encaps options...
      Send Auth=CHAP
      Recv PW=localpw
      Send PW=remotepw

```

3. Configure IP routing.

```

Route IP=Yes
IP options
  LAN Adrs=10.4.5.1/24
  RIP=Off

```

To make the MAX the default route for the site B Pipeline unit:

1. Open the Default IP Route profile in the site B Pipeline.
2. Specify the MAX unit at the far end of the WAN connection as the gateway address.

```

Ethernet
  Static Rtes
    Name=Default
    Active=Yes
    Dest=0.0.0/0
    Gateway=10.4.5.1
    Metric=1
    Preference=100
    Private=Yes

```

3. Close the IP Route profile.

Configuring a numbered interface

If you are not familiar with numbered interfaces, see [Numbered interfaces](#). In the following example, the MAX is a system-based router but supports a numbered interface for one of its connections. The arrow in [Figure 10-13](#) indicates the numbered interfaces for this connection:




Figure 10-13. Example numbered interface

The numbered interface addresses are:

- ! IF Adrs=10.5.6.7/24
- ! WAN Alias=10.7.8.9/24

An unnumbered interface is also shown in [Figure 10-13](#). The 10.1.2.3/32 connection uses a single system-based address for both the MAX itself and the dial-in user. To configure the numbered interface:

1. Open Ethernet > Mod Config > Ether Options and verify that the IP Adrs parameter is set correctly.

```

Ethernet
  Mod Config
    Ether options...
      IP Adrs=10.2.3.4/24
  
```

2. Close the Ethernet profile.
3. Open the Connection profile and configure the required parameters, then open the IP Options subprofile.
4. Specify the IP address of the remote device in the LAN Adrs parameter.

```

Ethernet
  Connections
    IP options...
      LAN Adrs=10.3.4.5/24
  
```

5. Specify the numbered interface address for the remote device in the WAN Alias parameter.

```

IP options...
  WAN Alias=10.7.8.9/24
  
```

6. Specify the numbered interface address for the 50 in the IF Adrs parameter.

```

IP options...
  IF Adrs=10.5.6.7/24
  
```

7. Close the Connection profile.

Configuring IP routes and preferences

The IP routing table contains routes that are configured (static routes) and routes that are learned dynamically from routing protocols such as RIP or OSPF. These are the parameters for configuring static routes:

```

Ethernet
  Static Rtes
    Name=route-name
    Active=Yes
    Dest=10.2.3.0/24
    Gateway=10.2.3.4
    Metric=2
    Preference=100
    Private=No
    Ospf=Cost=1
    ASE-type=Type1
    ASE=tag=c0000000

Ethernet
  Connections
    Route IP=Yes
    IP options...
      LAN Adrs=10.2.3.4/24
      WAN Alias=10.5.6.7/24
      IF Adrs=10.7.8.9/24
      Preference=100
      Metric=7
      DownPreference=120
      DownMetric=9
      Private=No

Ethernet
  Mod Config
    Ether options
      IP Adrs=10.2.3.1/24
      2nd Adrs=0.0.0.0/0
      RIP=Off

    Route Pref
      Static Preference=100
      Rip Preference=100
      RipAseType=Type2
      Rip Tag=c8000000
      OSPF Preference=10
      OSPF ASE Preference=150

```

Understanding the static route parameters

This section provides some background information on static routes. For more information on each parameter, see the *MAX Reference Guide*.

Route names

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

Activating a route

A route must be active to affect packet routing. An inactive route is ignored.

Route's destination address

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

Route's gateway address

The gateway-address parameter specifies the IP address of the router or interface to use to reach the target network.

Metrics, costs, and preferences

The metric parameter is a hop count for this route (a number between 1 to 15). When RIP was originally developed, the hop count was a number that showed how many routers needed to be crossed to reach the destination. For example, a destination with a hop count of 10 meant that to get a packet there requires crossing 10 routers. A route with a shorter hop count to a destination is more desirable than one with a larger hop count, since it most likely is a shorter, faster route.

The hop count can also be manually configured to give a route a "virtual" hop count. In this way you can manually configure which routes are more desirable than others in your environment. The higher the metric, the less likely that the MAX will use a route.

The cost parameter specifies the cost of an OSPF link. The cost is a configurable metric that can be used to take into account the speed of the link and other issues. The lower the cost, the more likely the interface will be used to forward data traffic. For details, see [Chapter 11, Configuring OSPF Routing](#).

The preference parameter specifies a route preference. Zero is the default for connected routes (such as the Ethernet). When choosing which route to use, the router first compares the preference values, preferring the lower number. If the preference values are equal, the router compares the metric values, using the route with the lower metric. The value of 255 means "Do not use this route." See [Route preferences and metrics](#).

Note: You can configure the DownMetric and DownPreference parameters to assign different metrics or preferences to routes on the basis of whether the route is in use or is down. You might want to direct the MAX to use active routes, if available, rather than choose routes that are down.

Tagging routes learned from RIP

The rip-tag field is *attached* to all routes learned from RIP in OSPF updates. The tag is a hexadecimal number that can be used by border routers to filter the record.

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

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

Making a route private

Private routes are used internally but are not advertised.

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

Routes for Connection profile interfaces

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

Note: You can configure the DownMetric and DownPreference parameters to assign different metrics or preferences to routes on the basis of whether the route is in use or is down. You might want to direct the MAX to use active routes, if available, rather than choose routes that are down.

A connected route for the Ethernet IP interface

The IP Adrs parameter specifies the MAX unit's IP address on the local Ethernet. The MAX creates a route for this address at system startup.

Static route preferences

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

RIP and OSPF preferences

Because OSPF typically involves a complex environment, its router configuration is described in a separate chapter. See [Chapter 11, Configuring OSPF Routing](#).

Tagging routes learned from RIP

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

Metrics for routes learned from RIP

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

Example static route configurations

For example Connection profile configurations, see [Configuring IP routing connections](#). Each of these results in a static route. For an example of the Ethernet profile configuration of the MAX unit's local IP interface, see [Configuring the MAX IP interface on a subnet](#).

Configuring the default route

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

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

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

1. Open the first IP Route profile (the route named Default) and activate it.

```
Ethernet
  Static Rtes
    Name=Default
    Active=Yes
    Dest=0.0.0.0/0
```

Note: The name of the first IP Route profile is always Default, and its destination is always 0.0.0.0 (you cannot change these values).

2. Specify the router to use for packets with unknown destinations; for example:

```
Gateway=10.9.8.10
```

3. Specify a metric for this route, the route's preference, and whether the route is private. For example:

```
Metric=1
Preference=100
Private=Yes
```

4. Close the IP Route profile.

Defining a static route to a remote subnet

If the connection does not enable RIP, the MAX does not learn about other networks or subnets that are reachable through the remote device, such as the remote network shown in [Figure 10-14](#).




Figure 10-14. Two-hop connection that requires a static route when RIP is off

To enable the MAX to route to site C without using RIP, you must configure an IP Route profile like this:

```
Ethernet
  Static Rtes
    Name=SITEBGW
    Active=Yes
    Dest=10.4.5.0/22
    Gateway=10.9.8.10
    Metric=2
    Preference=100
    Private=Yes
    Ospf=Cost=1
    ASE-type=Type1
    ASE=tag=c0000000
```

Example route preferences configuration

This example increases the preference value of RIP routes, instructing the router to use static routes first if one exists.

1. Open Ethernet > Mod Config > Route Pref.
2. Set Rip Preference to 150.

```
Ethernet
  Mod Config
    Route Pref
      Rip Preference=150
```

3. Close the Ethernet profile.

Configuring the MAX for dynamic route updates

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

These are the parameters that enable the MAX to receive updates from RIP or ICMP. (For information on OSPF updates, see [Chapter 11, Configuring OSPF Routing.](#))

```
Ethernet
```

```
Mod Config
  Ether options
    RIP=On
    Ignore Def Rt=Yes
  RIP Policy=Poison Rvrs
  RIP Summary=Yes
  ICMP Redirects=Accept

Ethernet
  Answer
    Session options...
    RIP=On

Ethernet
  Connections
    IP options...
    Private=No
    RIP=On
```

Understanding the dynamic routing parameters

This section provides some background information about the dynamic routing options.

RIP (Routing Information Protocol)

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

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

Ignoring the default route

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

RIP policy and RIP summary

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

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

The RIP Summary parameter specifies whether to summarize subnet information when advertising

routes. If the MAX summarizes RIP routes, it advertises a route to all the subnets in a network of the same class; for example, the route to 200.5.8.13/28 (a class C address subnetted to 28 bits) would be advertised as a route to 200.5.8.0. When the MAX does not summarize information, it advertises each route in its routing table "as-is;" in our example, the MAX advertises a route only to 200.5.8.13.

Ignoring ICMP Redirects

ICMP was designed to dynamically find the most efficient IP route to a destination. ICMP Redirect packets are one of the oldest route discovery methods on the Internet and one of the least secure, because it is possible to counterfeit ICMP Redirects and change the way a device routes packets.

Private routes

If you configure a profile with Private=Yes, the router will not disclose its route in response to queries from routing protocols.

Examples of RIP and ICMP configurations

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

1. Open Ethernet > Mod Config > Ether Options.
2. Configure the router to receive (but not send) RIP updates on Ethernet.

```
Ethernet
  Mod Config
    Ether options
      RIP=Recv-v2
```

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

3. Close the Ether Options subprofile, and set ICMP Redirects to Ignore.

```
ICMP Redirects=Ignore
```

4. Close the Ethernet profile.
5. Open Connections > IP Options, and configure the router to send (but not receive) RIP updates on this link.

```
Ethernet
  Connections
    IP options...
      RIP=Send-v2
```

Sending RIP on a WAN connection means that the remote devices will be able to access networks that are reachable via other local routers. However, the MAX does not receive information about networks that are reachable through the remote router.

6. Close the Connection profile.

Managing IP routes and connections

This section describes how to monitor TCP/IP/UDP and related information in the terminal server command-line interface. To invoke the terminal-server interface, select System > Sys Diag > Term Serv and press Enter.

Working with the IP routing table

The terminal-server IProute commands display the routing table and enable you to add or delete routes. The changes you make to the routing table using the IProute command last only until the MAX unit resets. To view the IProute commands:

```
ascend% iproute ?
```

```
iproute ?          Display help information
```

```
iproute add        iproute add <destination/size> <gateway> [ pref ] [ m
```

```
iproute delete    iproute delete <destination/size> <gateway> [ proto ]
```

```
iproute show      displays IP routes (same as "show ip routes" command)
```

Displaying the routing table

Note that the IProute Show command and the Show IP Routes command have identical output. To view the IP routing table:

```
ascend% iproute show
```

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

The columns in the table display the following information:

! Destination

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

! Gateway

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

! IF

The Interface column shows the name of the interface through which a packet addressed to this destination will be sent.

ie0 is the Ethernet interface

lo0 is the loopback interface

wanN specifies each of the active WAN interfaces

wanidle0 is the inactive interface (the special interface for any route whose WAN connection is down).

! Flg

The Flg column can contain the following flag values:

- " C (A directly connected route such as Ethernet)
- " I (ICMP Redirect dynamic route)
- " N (Placed in the table via SNMP MIB II)
- " O (A route learned from OSPF)
- " R (A route learned from RIP)
- " r (A RADIUS route)
- " S (A static route)
- " ? (A route of unknown origin, which indicates an error)
- " G (An indirect route via a gateway)
- " P (A private route)
- " T (A temporary route)
- " * (A hidden route that will not be used unless another better route to the same destination goes down)

! Pref

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

! Metric

The Metric column shows the RIP-style metric for the route, with a valid range of 0-16. Routes learned from OSPF show a RIP metric of 10. OSPF Cost infinity routes show a RIP metric of 16.

! Use

This is a count of the number of times the route was referenced since it was created. (Many of these references are internal, so this is not a count of the number of packets sent using this route.)

! Age

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

The first route in the default route (destination 0.0.0.0/0), which is pointing through the active Connection profile.

```
0.0.0.0/0          10.0.0.100      wan0      SG      1      1      0      20887
```

In this example, the IP Route profile for the default route specifies a Preference of 1, so this route is preferred over dynamically learned routes. The next route is specified in a Connection profile that is inactive.

```
10.207.76.0/24    10.207.76.1     wanidle0  SG      100     7      0      20887
```

The next route in the table is a static route that points through an inactive gateway:

```
10.207.77.0/24    10.207.76.1     wanidle0  SG      100     8      0      20887
```

The static route is followed by the loopback route:

```
127.0.0.1/32     -                lo0       CP      0      0      0      20887
```

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

The next route is specified in a Connection profile that is currently active:

```
10.0.0.0/24       10.0.0.100      wan0      SG      100     1      21387 20887
```

These are followed by the connection to the Ethernet interface. It is directly connected, with a Preference and Metric of zero.

```
10.1.2.0/24       -                ie0       C       0      0      19775 20887
```

The last two routes are a private loopback route, and a private route to the broadcast address:

```
10.1.2.1/32       -                lo0       CP      0      0      389   20887
```

```
255.255.255.255/32 -                ie0       CP      0      0      0     20887
```

The private loopback route is a host route with our Ethernet address. It is private, so it will not be advertised. The private route to the broadcast address is used in cases where the router will want to broadcast a packet but is otherwise unconfigured. It is typically used when trying to locate a server on a client machine to handle challenges for a token security card.

Adding an IP route

To add a static route to the MAX unit's routing table that will be lost when the MAX resets, use the IProute Add command in this format:

```
iproute add <destination> <gateway> [<metric>]
```

where <destination> is the destination network address, <gateway> is the IP address of the router that can forward packets to that network, and <metric> is the virtual hop count to the destination network (default 8). For example:

```
ascend% iproute add 10.1.2.0 10.0.0.3/24 1
```

The command shown immediately above adds a route to the 10.1.2.0 network and all of its subnets through the IP router located at 10.0.0.3/24. The metric to the route is 1 (it is one hop away).

If you try to add a route to a destination that already exists in the routing table, the MAX replaces the existing route, but only if the existing route has a higher metric. If you get the message "Warning: a better route appears to exist", the MAX rejected your attempt to add a route because the routing table already contained the same route with a lower metric. Note that RIP updates can change the metric for the route.

Deleting an IP route

To remove a route from the MAX unit's routing table, enter the IProute Delete command in this format:

```
iproute delete <destination> <gateway>
```

For example:

```
ascend% iproute delete 10.1.2.0 10.0.0.3/24
```

Note: RIP updates can add back any route you remove with IProute Delete. Also, the MAX restores all routes listed in the Static Route profile after a system reset.

Displaying route statistics

The Traceroute command is useful for locating slow routers or diagnosing IP routing problems. It traces the route an IP packet follows by launching UDP probe packets with a low TTL (Time-To-Live) value and then listening for an ICMP "time exceeded" reply from a router. Its syntax is:

```
traceroute [ -n ] [ -v ] [ -m max_ttl ] [ -p port ] [ -q nqueries ]  
[ -w waittime ] host [ datasize ]
```

All flags are optional. The only required parameter is the destination hostname or IP address.

! -n

Prints hop addresses numerically rather than symbolically and numerically (this eliminates a name server address-to-name lookup for each gateway found on the path).

! -v

Verbose output. Received ICMP packets other than Time Exceeded and ICMP Port Unreachable are listed.

! -m <max_ttl>

This sets the maximum time-to-live (maximum number of hops) used in outgoing probe packets. The default is 30 hops.

! -p <port>

Sets the base UDP port number used in probes. Traceroute hopes that nothing is listening on any of the UDP ports from the source to the destination host (so an ICMP Port Unreachable message will be returned to terminate the route tracing). If something is listening on a port in the default range, this option can be used to pick an unused port range. The default is 33434.

! -q <nqueries>

Sets the maximum number of queries for each hop. The default is 3.

! -w <waittime>

Sets the time to wait for a response to a query. The default is 3 seconds.

! host

The destination host by name or IP address.

! datasize

Sets the size of the data field of the UDP probe datagram sent by Traceroute. The default is 0. This results in a datagram size of 38 bytes (a UDP packet carrying no data).

For example, to trace the route to the host "techpubs":

```
ascend% traceroute techpubs

traceroute to techpubs (10.65.212.19), 30 hops MAX, 0 byte packets
 1  techpubs.eng.ascend.com (10.65.212.19)  0 ms  0 ms  0 ms
```

Probes start with a TTL of one and increase by one until of the following conditions occurs:

! The MAX receives an ICMP *port unreachable* message.

The UDP port in the probe packets is set to an unlikely value, such as 33434, because the target host is not intended to process the packets. A "port unreachable" message indicates that the packets reached the target host and were rejected.

- ! The TTL value reaches the maximum value.

By default, the maximum TTL is set to 30. You can specify a different TTL by using the `-m` option; for example:

```
ascend% traceroute -m 60 techpubs

traceroute to techpubs (10.65.212.19), 60 hops MAX, 0 byte packets
 1  techpubs.eng.abc.com (10.65.212.19)  0 ms  0 ms  0 ms
```

Three probes are sent at each TTL setting. The second line of command output shows the address of the router and round trip time of each probe. If the probe answers come from different gateways, the address of each responding system will be printed. If there is no response within a 3 second timeout interval, the command output is an asterisk. The following annotations may be included after the time field in a response:

- ! !H (Host reached.)
- ! !N (Network unreachable.)
- ! !P (Protocol unreachable.)
- ! !S (Source route failed. This may indicate a problem with the associated device.)
- ! !F (Fragmentation needed. This may indicate a problem with the associated device.)
- ! !h (Communication with the host is prohibited by filtering.)
- ! !n (Communication with the network is prohibited by filtering.)
- ! !c (Communication is otherwise prohibited by filtering.)
- ! !? (Indicates an ICMP sub-code. This should not occur.)
- ! !?? (Reply received with inappropriate type. This should not occur.)

Pinging other IP hosts

The terminal-server Ping command is useful for verifying that the transmission path is open between the MAX and another station. It sends an ICMP echo_request packet to the specified station. If the station receives the packet, it returns an ICMP echo_response packet. For example, to ping the *host techpubs*:

```
ascend% ping techpubs

PING techpubs (10.65.212.19): 56 data bytes
64 bytes from 10.65.212.19: icmp_seq=0 ttl=255 time=0 ms
64 bytes from 10.65.212.19: icmp_seq=3 ttl=255 time=0 ms
^C
--- techpubs ping statistics ---
```

```
2 packets transmitted, 2 packets received, 0% packet loss
```

```
round-trip min/avg/MAX = 0/0/0 ms
```

You can terminate the Ping exchange at any time by typing Ctrl-C. When you press Ctrl-C, the command reports the number of packets sent and received, the percentage of packet loss, duplicate or damaged echo_response packets (if any), and round-trip statistics. In some cases, round-trip times cannot be calculated.

During the Ping exchange, the MAX displays information about the packet exchange, including the TTL (Time-To-Live) of each ICMP echo_response packet.

Note: The maximum TTL for ICMP Ping is 255 and the maximum TTL for TCP is often 60 or lower, so you might be able to ping a host but not be able to run a TCP application (such as Telnet or FTP) to that station. If you Ping a host running a version of Berkeley UNIX before 4.3BSD-Tahoe, the TTL report is 255 minus the number of routers in the round-trip path. If you Ping a host running the current version of Berkeley UNIX, the TTL report is 255 minus the number of routers in the path from the remote system to the station performing the Ping.

The Ping command sends an ICMP mandatory echo_request datagram, which asks the remote station "Are you there?" If the echo_request reaches the remote station, the station sends back an ICMP echo_response datagram, which tells the sender "Yes, I am alive." This exchange verifies that the transmission path is open between the MAX and a remote station.

Configuring Finger support

You can configure the MAX to respond to Finger requests, as specified in RFC 1288-*The Finger User Information Protocol*.

To enable the MAX to respond to Finger requests:

1. Open the Ethernet > Mod Config.
2. Set Finger to Yes.
3. Exit and save the changes.

Displaying information

The following Show commands are useful for monitoring IP routing and related protocols:

```
show arp           Display the Arp Cache
show icmp          Display ICMP information
show if            Display Interface info. Type 'show if ?' for help.
show ip            Display IP information. Type 'show ip ?' for help.
show udp           Display UDP information. Type 'show udp ?' for help.
show tcp           Display TCP information. Type 'show tcp ?' for help.
```

show pools Display the assign address pools.

Displaying the ARP cache

To view the ARP cache:

```
ascend% show arp
```

entry	typ	ip address	ether addr	if	rtr	pkt	insert
0	DYN	10.65.212.199	00C07B605C07	0	0	0	857783
1	DYN	10.65.212.91	0080C7C4CB80	0	0	0	857866
2	DYN	10.65.212.22	080020792B4C	0	0	0	857937
3	DYN	10.65.212.3	0000813DF048	0	0	0	857566
4	DYN	10.65.212.250	0020AFF80F1D	0	0	0	857883
5	DYN	10.65.212.16	0020AFEC0AFB	0	0	0	857861
6	DYN	10.65.212.227	00C07B5F14B6	0	0	0	857479
7	DYN	10.65.212.36	00C07B5E9AA5	0	0	0	857602
8	DYN	10.65.212.71	0080C730041F	0	0	0	857721
9	DYN	10.65.212.5	0003C6010512	0	0	0	857602
10	DYN	10.65.212.241	0080C72ED212	0	0	0	857781
11	DYN	10.65.212.120	0080C7152582	0	0	0	857604
12	DYN	10.65.212.156	0080A30ECE6D	0	0	0	857901
13	DYN	10.65.212.100	00C07B60E28D	0	0	0	857934
14	DYN	10.65.212.1	00000C065D27	0	0	0	857854
15	DYN	10.65.212.102	08000716C449	0	0	0	857724
16	DYN	10.65.212.33	00A024AA0283	0	0	0	857699
17	DYN	10.65.212.96	0080C7301792	0	0	0	857757
18	DYN	10.65.212.121	0080C79BF681	0	0	0	857848
19	DYN	10.65.212.89	00A024A9FB99	0	0	0	857790
20	DYN	10.65.212.26	00A024A8122C	0	0	0	857861
21	DYN	10.65.212.6	0800207956A2	0	0	0	857918
22	DYN	10.65.212.191	0080C75BE778	0	0	0	857918
23	DYN	10.65.212.116	0080C72F66CC	0	0	0	857416
24	DYN	10.65.212.87	0000813606A0	0	0	0	857666
25	DYN	10.65.212.235	00C07B76D119	0	0	0	857708
26	DYN	10.65.212.19	08002075806B	0	0	0	857929

The ARP table displays this information:

- ! entry: A unique identifier for each ARP table entry.
- ! typ: How the address was learned, dynamically (DYN) or statically (STAT).
- ! ip address: The address contained in ARP requests.
- ! ether addr: The MAC address of the host with that IP address.
- ! if: The interface on which the MAX received the ARP request.
- ! rtr: The next-hop router on the specified interface.

Displaying ICMP packet statistics

To view the number of ICMP packets received intact, received with errors, and transmitted:

```
ascend% show icmp

3857661 packet received.
20 packets received with errors.
   Input histogram: 15070
2758129 packets transmitted.
0 packets transmitted due to lack of resources.
   Output histogram: 15218
```

The Input and Output histograms show the number of ICMP packets received and transmitted in each category.

Displaying interface statistics

To see the supported commands:

```
ascend% show if ?

show if ?          Display help information
show if stats      Display Interface Statistics
show if totals     Display Interface Total counts
```

To display the status and packet count of each active WAN link as well as local and loopback interfaces:

```
ascend% show if stats
```

Interface	Name	Status	Type	Speed	MTU	InPackets	Outpacket
ie0	ethernet	Up	6	10000000	1500	107385	85384
wan0		Down	1	0	1500	0	0
wan1		Down	1	0	1500	0	0
wan2		Down	1	0	1500	0	0

```

wanidle0          Up    6    10000000    1500    0    0
lo0      loopback  Up   24    10000000    1500    0    0

```

The output contains these fields:

- ! Interface: The interface name (see [Chapter 10, Configuring IP Routing.](#))
- ! Name: The name of the profile or a text name for the interface
- ! Status: Up (the interface is functional) or Down.
- ! Type: The type of application being used on the interface, as specified in RFC 1213 (MIB-2). For example, 23 indicates PPP and 28 indicates SLIP.
- ! Speed: The data rate in bits per second.
- ! MTU: The maximum packet size allowed on the interface. MTU stands for Maximum Transmission Unit.
- ! InPackets: The number of packets the interface has received.
- ! OutPackets: The number of packets the interface has transmitted.

To display the packet count at each interface broken down by type of packet:

```
ascend% show if totals
```

```

Name  --Octets----Ucast-- -NonUcast- Discard -Error- Unknown -Same IF-
ie0   i:    7813606    85121      22383      0      0      0      0
      o:   101529978    85306       149        0      0      0      0
wan0  i:         0         0         0         0      0      0      0
      o:         0         0         0         0      0      0      0
wan1  i:         0         0         0         0      0      0      0
      o:         0         0         0         0      0      0      0
wan2  i:         0         0         0         0      0      0      0
      o:         0         0         0         0      0      0      0
wanidle0 i:         0         0         0         0      0      0      0
        o:         0         0         0         0      0      0      0
lo0   i:         0         0         0         0      0      0      0
      o:         0         0         0         0      0      0      0

```

The output contains these fields:

- ! Name: The interface name (see [Chapter 10, Configuring IP Routing.](#))

- ! Octets: The total number of bytes processed by the interface.
- ! Ucast: Packets with a unicast destination address.
- ! NonUcast: Packets with a multicast address or a broadcast address.
- ! Discard: The number of packets that the interface could not process.
- ! Error: The number of packets with CRC errors, header errors, or collisions.
- ! Unknown: The number of packets the MAX forwarded across all bridged interfaces because of unknown or unlearned destinations.
- ! Same IF: The number of bridged packets whose destination is the same as the source.

Displaying IP statistics and addresses

To see the supported commands:

```
ascend% show ip ?  
  
show ip ?          Display help information  
show ip stats      Display IP Statistics  
show ip address    Display IP Address Assignments  
show ip routes     Display IP Routes
```

Note: For information on the Show IP Routes command, see [Working with the IP routing table](#).

To display statistics on IP activity, including the number of IP packets the MAX has received and transmitted:

```
ascend% show ip stats  
  
107408 packets received.  
    0 packets received with header errors.  
    0 packets received with address errors.  
    0 packets forwarded.  
    0 packets received with unknown protocols.  
    0 inbound packets discarded.  
  
107408 packets delivered to upper layers.  
    85421 transmit requests.  
    0 discarded transmit packets.  
    1 outbound packets with no route.
```

```

0 reassembly timeouts.
0 reassemblies required.
0 reassemblies that went OK.
0 reassemblies that Failed.
0 packets fragmented OK.
0 fragmentations that failed.
0 fragment packets created.
0 route discards due to lack of memory.
64 default ttl.

```

To view IP interface address information:

```
ascend% show ip address
```

Interface	IP Address	Dest Address	Netmask	MTU	Status
ie0	10.2.3.4	N/A	255.255.255.224	1500	Up
wan0	0.0.0.0	N/A	0.0.0.0	1500	Down
wan1	13.1.2.0	13.1.2.128	255.255.255.248	1500	Down
wan2	0.0.0.0	N/A	0.0.0.0	1500	Down
wan3	0.0.0.0	N/A	0.0.0.0	1500	Down
lo0	127.0.0.1	N/A	255.255.255.255	1500	Up
rj0	127.0.0.2	N/A	255.255.255.255	1500	Up
bh0	127.0.0.3	N/A	255.255.255.255	1500	Up

Displaying UDP statistics and listen table

To see the supported commands:

```
ascend% show udp ?
```

```

show udp ?          Display help information
show udp stats      Display UDP Statistics
show udp listen     Display UDP Listen Table

```

To display the number of UDP packets received and transmitted:

```
ascend% show udp stats
```

```

22386 packets received.
    0 packets received with no ports.
    0 packets received with errors.

```

0 packets dropped

9 packets transmitted.

In addition to the socket number, UDP port number and the number of packets queued for each UDP port on which the MAX is currently listening, the show udp listen command now shows these additional parameters:

- ! InQMax - The maximum number of queued UDP packets on the socket (See Queue Depth and Rip Queue Depth parameters.)
- ! InQLen - The current number of queued packets on the socket
- ! InQDrops - The number of packets discarded because it would cause InQLen to exceed InQMax
- ! Total Rx - The total number of packets received on the socket, including InQDrops

An example follows:

```
ascend% show udp listen
```

```
udp:
```

Socket	Local Port	InQLen	InQMax	InQDrops	Total Rx	
0	1023	0	1	0	0	
1	520	0	50	0	532	
2	7	0	32	0	0	
3	123	0	32	0	0	
4	1022	0	128	0	0	
5		161	0	64	0	0

Displaying TCP statistics and connections

To see the supported commands:

```
ascend% show tcp ?
```

```
show tcp ?          Display help information
```

```
show tcp stats      Display TCP Statistics
```

```
show tcp connection Display TCP Connection Table
```

To display the number of TCP packets received and transmitted:

```
ascend% show tcp stats
```

```
0 active opens.
```

```
11 passive opens.
```

```

1 connect attempts failed.
1 connections were reset.
3 connections currently established.
85262 segments received.
85598 segments transmitted.
559 segments re-transmitted.

```

An active open is a TCP session that the MAX initiated, and a passive open is a TCP session that the MAX did not initiate.

To display current TCP sessions:

```
ascend% show tcp connection
```

Socket	Local	Remote	State
0	*.23	*.*	LISTEN
1	10.2.3.23	15.5.248.121.15003	ESTABLISHED

Displaying address pool status

To view the status of the MAX unit's IP address pool:

```
ascend% show pools
```

Pool #	Base	Count	InUse
1	10.98.1.2	55	27
2	10.5.6.1	128	0
Number of remaining allocated addresses:			156

[HOME](#) [CONTENTS](#) [PREVIOUS](#) [NEXT](#) [INDEX](#)

techpubs@eng.ascend.com

Copyright © 1998, Ascend Communications, Inc. All rights reserved.