

# Configuring IP

---

The Internet Protocol (IP) is a packet-based protocol used to exchange data over computer networks. IP handles addressing, fragmentation, reassembly, and protocol demultiplexing. It is the foundation on which all other Internet protocols, collectively referred to as the Internet Protocol suite, are built. IP is a network-layer protocol that contains addressing information and some control information that allows data packets to be routed.

The Transmission Control Protocol (TCP) is built upon the IP layer. TCP is a connection-oriented protocol that specifies the format of data and acknowledgments used in the transfer of data. TCP also specifies the procedures that the computers use to ensure that the data arrives correctly. TCP allows multiple applications on a system to communicate concurrently because it handles all demultiplexing of the incoming traffic among the application programs.

This chapter describes how to configure the IP protocol. For a complete description of the commands in this chapter, refer to the “IP Commands” chapter of the *Router Products Command Reference* publication. For information on configuring the various IP routing protocols, refer to the “Configuring IP Routing Protocols” chapter of this manual. For historical background and a technical overview of IP and other protocols, see the *Internetworking Technology Overview* publication.

## Cisco’s Implementation of IP

Cisco’s implementation of TCP/IP provides most of the major services contained in the various protocol specifications. Cisco routers also provide the TCP and UDP services called Echo and Discard, which are described in RFCs 862 and 863, respectively.

Cisco supports both TCP and UDP at the transport layer, for maximum flexibility in services. Cisco also supports all standards for IP broadcasts.

## IP Configuration Task List

A number of tasks are associated with configuring IP. A basic and required task for configuring IP is to assign IP addresses to network interfaces. Doing so enables the interfaces and allows communication with hosts on those interfaces using IP. Associated with this task are decisions about subnetting and masking the IP addresses.

The following list summarizes these tasks and features:

- Assign IP addresses to network interfaces (page 15-2)
- Configure IP addressing options (page 15-5)
- Disable IP routing (page 15-10)
- Configure a routing process (page 15-13)
- Configure broadcast packet handling (page 15-13)
- Configure IP services (page 15-17)
- Filter IP packets (page 15-21)
- Configure IP security options (page 15-23)
- Configure IP accounting (page 15-25)
- Configure performance parameters (page 15-26)
- Configure IP over WANs (page 15-29)
- Monitor and maintain the IP network (page 15-29)

Each task is described in a following section. Remember that not all of these tasks are required. The tasks you need to perform will depend on your network and your needs.

At the end of this chapter (page 15-31) are examples illustrating how you might configure your network using IP.

## Assign IP Addresses to Network Interfaces

IP addresses identify locations to which IP datagrams can be sent. See the *Internetworking Technology Overview* publication for detailed information on IP addresses.

To assign an Internet address and a network mask to a network interface on the router, perform the following task in interface configuration mode:

---

Task	Command
Set an IP address for an interface.	<b>ip address</b> <i>IP-address mask</i>

---

A mask identifies the bits that denote the network number in an IP address. When you use the mask to subnet a network, the mask is then referred to as a *subnet mask*. Subnets are described in the *Internetworking Technology Overview* publication.

---

**Note** We only support network masks that use contiguous bits that are flush left against the network field.

---

## Assign Multiple IP Addresses to Network Interfaces

The software supports multiple IP addresses per interface. You can specify an unlimited number of secondary addresses. Secondary IP addresses can be used in a variety of situations. The following are the most common applications:

- There may not be enough host addresses for a particular network segment. For example, your subnetting allows up to 254 hosts per logical subnet, but on one physical subnet you need to have 300 host addresses. Using secondary IP addresses on the routers allows you to have two logical subnets using one physical subnet.
- Many older networks were built using Level 2 bridges, and were not subnetted. The judicious use of secondary addresses can aid in the transition to a subnetted, router-based network. Routers on an older, bridged segment can easily be made aware that there are many subnets on that segment.
- Two subnets of a single network might otherwise be separated by another network. You can create a single network from subnets that are physically separated by another network by using a secondary address. In these instances, the first network is *extended*, or layered on top of the second network. Note that a subnet cannot appear on more than one active interface of the router at a time.

---

**Note** If any router on a network segment uses a secondary address, all other routers on that same segment must also use a secondary address from the same network or subnet.

---

To assign multiple IP addresses to network interfaces, perform the following task in interface configuration mode:

Task	Command
Assign multiple IP addresses to network interfaces.	<b>ip address <i>IP-address mask secondary</i></b>

See the “IP Configuration Examples” section at the end of the chapter for an example of creating a network from separated subnets.

## Allowable Internet Addresses

Some Internet addresses are reserved for special uses and cannot be used for host, subnet, or network addresses. Table 1-1 lists ranges of Internet addresses and shows which addresses are reserved and which are available for use.

**Table 1-1 Reserved and Available Internet Addresses**

Class	Address or Range	Status
A	0.0.0.0	Reserved
	1.0.0.0 through 126.0.0.0	Available
	127.0.0.0	Reserved
B	128.0.0.0	Reserved
	128.1.0.0 through 191.254.0.0	Available
	191.255.0.0	Reserved
C	192.0.0.0	Reserved
	192.0.1.0 through 223.255.254	Available
	223.255.255.0	Reserved
D	224.0.0.0 through 239.255.255.255	Multicast group addresses
E	240.0.0.0 through 255.255.255.254	Reserved
	255.255.255.255	Broadcast

The official description of Internet addresses is found in RFC 1166, “Internet Numbers.”

To receive an assigned network number, contact your IP service provider.

## Enable Use of Subnet Zero

Subnetting with a subnet address of zero is illegal and strongly discouraged (as stated in RFC 791) because of the confusion that can arise between a network and a subnet that have the same addresses. For example, if network 131.108.0.0 is subnetted as 255.255.255.0, subnet zero would be written as 131.108.0.0—which is identical to the network address.

There is a way to use the all zeros and all ones subnet (131.108.255.0), even though it is discouraged. Configuring interfaces for the all ones subnet is explicitly allowed. However, if you need the entire subnet space for your IP address, perform the following task in global configuration mode to enable subnet zero:

Task	Command
Enable the use of subnet zero for interface addresses and routing updates.	<b>ip subnet-zero</b>

## Enable IP Processing on a Serial Interface

You might want to enable IP processing on a serial or tunnel interface without assigning an explicit IP address to the interface. Whenever the unnumbered interface generates a packet (for example, for a routing update), it uses the address of the interface you specified as the source address of the IP packet. It also uses the specified interface address in determining which routing processes are sending updates over the unnumbered interface. Restrictions are as follows:

- Point-to-point Frame Relay interfaces can be unnumbered. It is not possible to use unnumbered interfaces with HDLC, PPP, LAPB, X.25, or SMDS encapsulations, or with SLIP and tunnel interfaces.
- You cannot use the **ping EXEC** command to determine whether the interface is up, because the interface has no IP address. Simple Network Management Protocol (SNMP) can be used to remotely monitor interface status.
- You cannot netboot a runnable image over an unnumbered serial interface.
- You cannot support IP security options on an unnumbered interface.

If you are configuring IS-IS across a serial line, you should configure the serial interfaces as unnumbered. This allows you to conform with RFC 1195, which states that IP addresses are not required on each interface.

---

**Note** Using an unnumbered serial line between different major networks (majornets) requires special care. If at each end of the link there are different majornets assigned to the interfaces you specified as unnumbered, any routing protocols running across the serial line cannot advertise subnet information.

---

To enable IP processing on an unnumbered serial interface, perform the following task in interface configuration mode:

Task	Command
Enable IP processing on a serial or tunnel interface without assigning an explicit IP address to the interface.	<b>ip unnumbered</b> <i>interface-name</i>

The interface you specify must be the name of another interface in the network that has an IP address, not another unnumbered interface.

The interface you specify also must be enabled (listed as “up” in the **show interfaces** command display).

An example of how to configure serial interfaces can be found in the “IP Configuration Examples” section at the end of the chapter.

## Configure IP Addressing Options

Our IP implementation allows you to control interface-specific handling of IP addresses by facilitating address resolution, name services, and other functions. The following tasks are described in this section:

- Establish address resolution
- Map host names to IP addresses
- Configure HP Probe Proxy name requests

The following sections provide additional information and describe these tasks.

## Establish Address Resolution

A device in the Internet can have both a local address, which uniquely identifies the device on its local segment or LAN, and a network address, which identifies the network the device belongs to. The local address is more properly known as a data link address, because it is contained in the data link layer (Layer 2 of the OSI model) part of the packet header and is read by data link devices (bridges and all device interfaces, for example). The more technically inclined will refer to local addresses as MAC addresses, because the Media Access Control (MAC) sublayer within the data link layer processes addresses for the layer.

To communicate with a device on Ethernet, for example, the router first must determine the 48-bit MAC or local data link address of that device. The process of determining the local data link address from an Internet address is called *address resolution*. The process of determining the Internet address from a local data link address is called *reverse address resolution*. The router uses three forms of address resolution: Address Resolution Protocol (ARP), proxy ARP, and Probe (which is similar to ARP). The router also uses the Reverse Address Resolution Protocol (RARP). The ARP, proxy ARP, and RARP protocols are defined in RFCs 826, 1027, and 903, respectively. Probe is a protocol developed by the Hewlett-Packard Company for use on IEEE-802.3 networks.

The *Address Resolution Protocol (ARP)* is used to associate IP addresses with media or MAC addresses. Taking an IP address as input, ARP determines the associated media address. Once a media or MAC address is determined, the IP address/media address association is stored in an ARP cache for rapid retrieval. Then the IP datagram is encapsulated in a link-layer frame and sent over the network. Encapsulation of IP datagrams and ARP requests and replies on IEEE 802 networks other than Ethernet is specified by the *Subnetwork Access Protocol (SNAP)*.

The *Reverse Address Resolution Protocol (RARP)* works the same way as ARP, except that the RARP Request packet requests an Internet address instead of a local data link address. Use of RARP requires a RARP server on the same network segment as the router interface. RARP often is used by diskless nodes that do not know their IP addresses when they boot. Our routers attempt to use RARP if they do not know the IP address of an interface at startup. Also, the routers are able to act as RARP servers by responding to RARP requests that they are able to answer. See the “Loading System Images, Microcode Images, and Configuration Files” chapter to learn how to configure a router as a RARP server.

Perform the following tasks to set address resolution on the router:

- Define a static ARP cache, as necessary
- Set ARP encapsulations
- Set Proxy ARP

The procedures for performing these tasks are described in the following sections.

### Define a Static ARP Cache

ARP and other address resolution protocols provide a dynamic mapping between IP addresses and media addresses. Because most hosts support dynamic address resolution, you generally do not need to specify static ARP cache entries. If you do need to define them, you can do so globally. Doing this task installs a permanent entry in the ARP cache. The router uses this entry to translate 32-bit IP addresses into 48-bit hardware addresses.

Optionally, you can specify that the router respond to ARP requests as if it were the owner of the specified IP address, and you also have the option of specifying an ARP entry timeout period when you define ARP entries.

The following two tables list the tasks to provide dynamic mapping between IP addresses and media address.

Perform either of the following tasks in global configuration mode:

Task	Command
Globally associate an Internet address with a media (hardware) address in the ARP cache.	<b>arp</b> <i>internet-address hardware-address type</i>
Specify that the router respond to ARP requests as if it were the owner of the specified IP address.	<b>arp</b> <i>internet-address hardware-address type alias</i>

Perform the following task in interface configuration mode:

Task	Command
Set the length of time an ARP cache entry will stay in the cache.	<b>arp</b> <i>timeout seconds</i>

To display the type of ARP being used on a particular interface and also display the ARP timeout value, use the **show interfaces EXEC** command. Use the **show arp EXEC** command to examine the contents of the ARP cache. Use the **show ip arp EXEC** command to show IP entries. To remove all nonstatic entries from the ARP cache, use the privileged EXEC command **clear arp-cache**.

## Set ARP Encapsulations

By default, standard Ethernet-style ARP encapsulation (represented by the **arpa** keyword) is enabled on the IP interface. You can change this encapsulation method to SNAP or HP Probe, as required by your network, to control the interface-specific handling of IP address resolution into 48-bit Ethernet hardware addresses.

When you set HP Probe encapsulation, the router uses the Probe protocol whenever it attempts to resolve an IEEE-802.3 or Ethernet local data link address. The subset of Probe that performs address resolution is called *Virtual Address Request and Reply*. Using Probe, the router can communicate transparently with Hewlett-Packard IEEE-802.3 hosts that use this type of data encapsulation. You must explicitly configure all interfaces for Probe that will use Probe.

To specify the ARP encapsulation type, perform the following task in interface configuration mode:

Task	Command
Specify one of three ARP encapsulation methods for a specified interface.	<b>arp</b> { <b>arpa</b>   <b>probe</b>   <b>snap</b> }

## Disable Proxy ARP

The router uses proxy ARP, as defined in RFC 1027, to help hosts with no knowledge of routing determine the media addresses of hosts on other networks or subnets. For example, if the router receives an ARP request for a host that is not on the same network as the ARP request sender, and

if the router has the best route to that host, then the router sends an ARP reply packet giving its own local data link address. The host that sent the ARP request then sends its packets to the router, which forwards them to the intended host. Proxy ARP is enabled by default.

To disable proxy ARP, perform the following task in interface configuration mode, as necessary, for your network:

Task	Command
Disable proxy ARP on the interface.	<b>no ip proxy-arp</b>

## Map Host Names to IP Addresses

Each unique IP address can have a host name associated with it. The router maintains a cache of host name-to-address mappings for use by the EXEC **connect**, **telnet**, **ping** and related Telnet support operations. This cache speeds the process of converting names to addresses.

IP defines a naming scheme that allows a device to be identified by its location in the Internet. This is a hierarchical naming scheme that provides for *domains*. Domain names are pieced together with periods (.) as the delimiting characters. For example, Cisco Systems is a commercial organization that the Internet identifies by a *com* domain name, so its domain name is *cisco.com*. A specific device in this domain, the FTP system for example, is identified as *ftp.cisco.com*.

To keep track of domain names, IP has defined the concept of a *name server* whose job it is to hold a cache, or database, of names mapped to IP addresses. To map domain names to IP addresses, you must first identify the host names, then specify a name server, and enable the Domain Name System (DNS), the Internet’s global naming scheme that uniquely identifies network devices. You do these by performing the following tasks:

- Map IP addresses to host names
- Specify the domain name
- Specify a name server
- Disable the DNS
  - Use the DNS to discover CLNS addresses

The following sections describe these tasks.

### Map IP Addresses to Host Names

The router maintains a table of host names and their corresponding addresses, also called a *host name-to-address mapping*. Higher-layer protocols such as Telnet use host names to identify network devices (hosts). The router and other network devices must be able to associate host names with IP addresses to communicate with other IP devices. Host names and IP addresses can be associated with one another through static or dynamic means.

Manually assigning host names to addresses is useful when dynamic mapping is not available.

To assign host names to addresses, perform the following task in global configuration mode:

Task	Command
Statically associate host names with IP addresses.	<b>ip host name</b> [ <i>TCP-port-number</i> ] <i>address1</i> [ <i>address2...address8</i> ]



## Specify the Domain Name

You can specify a default domain name that the router software will use to complete domain name requests. You can specify either a single domain name or a list of domain names. Any IP host name that does not contain a domain name will have the domain name you specify appended to it before being added to the host table.

To specify a domain name or names, perform either of the following tasks in global configuration mode:

Task	Command
Define a default domain name that the router will use to complete unqualified host names.	<b>ip domain-name</b> <i>name</i>
or	
Define a list of default domain names to complete unqualified host names.	<b>ip domain-list</b> <i>name</i>

See the “IP Configuration Examples” section at the end of this chapter for an example of establishing IP domains.

## Specify a Name Server

To specify one or more hosts (up to six) that can function as a name server to supply name information for the DNS, perform the following task in global configuration mode:

Task	Command
Specify one or more hosts that supply name information.	<b>ip name-server</b> <i>server-address1</i> [[ <i>server-address2</i> ]... <i>server-address6</i> ]

## Disable the DNS

If your network devices require connectivity with devices in networks for which you do not control name assignment, you can assign device names that uniquely identify your devices within the entire internetwork. The Internet’s global naming scheme, the *Domain Name System (DNS)*, accomplishes this task. This service is enabled by default.

To disable the DNS, perform the following task in global configuration mode:

Task	Command
Disable DNS-based host name-to-address translation.	<b>no ip domain-lookup</b>

See the “IP Configuration Examples” section at the end of this chapter for an example of enabling the DNS.

### Use the DNS to Discover ISO CLNS Addresses

If your router has both IP and ISO CLNS enabled and you want to use ISO CLNS NSAP (Network Service Access Point) addresses, you can use the DNS to query these addresses, as documented in RFC 1348. This feature is enabled by default.

To disable DNS queries for ISO CLNS addresses, perform the following task in global configuration mode:

Task	Command
Disable DNS queries for ISO CLNS addresses.	<b>no ip domain-lookup nsap</b>

### Configure HP Probe Proxy Name Requests

HP Probe Proxy support allows a router to respond to HP Probe Proxy name requests. These requests are typically used at sites that have Hewlett-Packard (HP) equipment and are already using HP Probe. Tasks associated with HP Probe Proxy are shown in the following two tables.

To configure HP Probe Proxy, perform the following task in interface configuration mode:

Task	Command
Allow the router to respond to HP Probe Proxy name requests.	<b>ip probe proxy</b>

Perform the following task in global configuration mode:

Task	Command
Enter the host name of an HP host (for which the router is acting as a proxy) into the host table.	<b>ip hp-host</b> <i>hostname ip-address</i>

See the “IP Configuration Examples” section at the end of this chapter for an example of configuring HP hosts on a network segment.

## Disable IP Routing

Every router ships with IP routing automatically enabled. If you choose to set up the router to bridge rather than route IP datagrams, you must disable IP routing. To disable IP routing, perform the following task in global configuration mode:

Task	Command
Disable IP routing.	<b>no ip routing</b>

When IP routing is disabled, the router will act as an IP end host for IP packets destined for or sourced by it, whether or not bridging is enabled for those IP packets not destined for the router. To reenabling IP routing, use the **ip routing** command.

## Routing Assistance when IP Routing Is Disabled

The router software provides three methods by which the router can learn about routes to other networks when IP routing is disabled and the router is acting as an IP host:

- Proxy ARP (Address Resolution Protocol)
- A default gateway (also known as default router)
- The Router Discovery mechanism

When IP routing is disabled, the default gateway feature and the router discovery client are enabled, and Proxy ARP is disabled. When IP routing is enabled, the default gateway feature is disabled and you can configure Proxy ARP and the router discovery servers.

### Proxy ARP

The most common method of learning about other routes is by using proxy ARP. Proxy ARP, defined in RFC 1027, enables an Ethernet host with no knowledge of routing to communicate with hosts on other networks or subnets. Such a host assumes that all hosts are on the same local Ethernet and that it can use ARP to determine their hardware addresses.

Under proxy ARP, if a router receives an ARP Request for a host that is not on the same network as the ARP Request sender, the router evaluates whether it has the best route to that host. If the router does have the best route, it sends an ARP Reply packet giving its own Ethernet hardware address. The host that sent the ARP Request then sends its packets to the router, which forwards them to the intended host. The software treats all networks as if they are local and performs ARP requests for every IP address. This feature is enabled by default.

Proxy ARP works as long as other routers support it. Many other routers, especially host-based routing software, do not support it.

### A Default Gateway

Another method for locating routes is to define a default router (or gateway). The software sends all nonlocal packets to this router, which either routes them appropriately or sends an ICMP redirect message back to the router, telling it of a better route. The ICMP redirect message indicates which local router the host should use. The software caches the redirect messages and routes each packet thereafter as efficiently as possible. The limitations of this method are that there is no means of detecting when the default router has crashed or is unavailable and no method of picking another router if one of these events should occur.

To set up a default gateway for a host, perform the following task in global configuration mode:

Task	Command
Set up a default gateway (router).	<b>ip default-gateway</b> <i>address</i>

To display the address of the default gateway, use the **show ip redirects EXEC** command.

### The Router Discovery Mechanism

The router software provides a third method, called *router discovery*, by which the router can dynamically learn about routes to other networks using the Gateway Discovery Protocol (GDP) or the ICMP Router Discovery Protocol (IRDP) for detecting routers. The software is also capable of

wiretapping RIP and IGRP routing updates and inferring the location of routers from those updates. The server/client implementation of router discovery does not actually examine or store the full routing tables sent by routers, it merely keeps track of which systems are sending such data.

This mechanism supports the following protocols:

- Gateway Discovery Protocol (GDP)
- ICMP Router Discovery Protocol (IRDP)
- Routing Information Protocol (RIP)
- Interior Gateway Routing Protocol (IGRP)

You can configure these protocols in any combination. When possible, we recommend that you use GDP or IRDP, because they allow each router to specify *both* a priority and the time after which a router should be assumed down if no further packets are received. Routers discovered using IGRP are assigned an arbitrary priority of 60. Routers discovered through RIP are assigned a priority of 50. For IGRP and RIP, the software attempts to measure the time between updates and will assume that the router is down if no updates are received for 2.5 times that interval.

Each router discovered becomes a candidate for the default router. The list of candidates is scanned and a new highest-priority router is selected when any of the following events occur:

- When a higher-priority router is discovered (the list of routers is polled at 5-minute intervals).
- When the current default router is declared down.
- When a TCP connection is about to time out because of excessive retransmissions. In this case, the server flushes the ARP cache and the ICMP redirect cache and picks a new default router in an attempt to find a successful route to the destination.

### Detect Routers Using the GDP Protocol

To configure the router discovery feature using the GDP routing protocol, perform the following task in interface configuration mode:

Task	Command
Use the GDP protocol to configure router discovery.	<b>ip gdp gdp</b>

Use the **no ip gdp gdp** command to turn the router discovery feature off.

### Detect Routers Using the IRDP Protocol

To configure the router discovery feature using the IRDP routing protocol, perform the following task in interface configuration mode:

Task	Command
Use the IRDP protocol to configure router discovery.	<b>ip gdp irdp</b>

Use the **no ip gdp irdp** command to turn the router discovery feature off.

### Detect Routers Using the RIP Protocol

To configure the router discovery feature using the RIP routing protocol, perform the following task in interface configuration mode:

Task	Command
Use the RIP protocol to configure router discovery.	<b>ip gdp rip</b>

Use the **no ip gdp rip** command to turn the router discovery feature off.

### Detect Routers Using the IGRP Protocol

To configure the router discovery feature using the IGRP routing protocol, perform the following task in interface configuration mode:

Task	Command
Use the IGRP protocol to configure router discovery.	<b>ip gdp igrp</b>

Use the **no ip gdp igrp** command to turn the router discovery feature off.

## Configure a Routing Process

At this point in the configuration process, you may choose to configure one or more of the many routing protocols that are available based on your individual network needs. Routing protocols provide topology information of an internetwork. Refer to the “Configuring IP Routing Protocols” chapter for the tasks involved in configuring IP routing protocols. If you want to continue to perform basic IP configuration tasks, continue reading the following sections.

## Configure Broadcast Packet Handling

A *broadcast* is a data packet destined for all hosts on a particular physical network. Network hosts recognize broadcasts by special addresses. Broadcasts are heavily used by some protocols, including several important Internet protocols. Control of broadcast messages is an essential part of the IP network administrator’s job.

Our routers support two kinds of broadcasting: *directed broadcasting* and *flooding*. A directed broadcast is a packet sent to a specific network or series of networks, while a flooded broadcast packet is sent to every network. A directed broadcast address includes the network or subnet fields.

Several early IP implementations do not use the current broadcast address standard. Instead, they use the old standard, which calls for all zeros instead of all ones to indicate broadcast addresses. Many of these implementations do not recognize an all-ones broadcast address and fail to respond to the broadcast correctly. Others forward all-ones broadcasts, which causes a serious network overload known as a *broadcast storm*. Implementations that exhibit these problems include systems based on versions of BSD UNIX prior to Version 4.3.

Routers provide some protection from broadcast storms by limiting their extent to the local cable. Bridges (including intelligent bridges), because they are Layer 2 devices, forward broadcasts to all network segments, thus propagating all broadcast storms.

The best solution to the broadcast storm problem is to use a single broadcast address scheme on a network. Most modern IP implementations allow the network manager to set the address to be used as the broadcast address. Many implementations, including the one on our router, can accept and interpret all possible forms of broadcast addresses.

For detailed discussions of broadcast issues in general, see RFC 919, "Broadcasting Internet Datagrams," and RFC 922, "Broadcasting Internet Datagrams in the Presence of Subnets." The router support for Internet broadcasts generally complies with RFC 919 and RFC 922; however, it does not support multisubnet broadcasts as defined in RFC 922.

The current broadcast address standard provides specific addressing schemes for forwarding broadcasts. Perform the following tasks to enable these schemes:

- Enable directed broadcast-to-physical broadcast translation
- Forward UDP broadcast packets and protocols
- Establish an Internet broadcast address
- Flood IP broadcasts

Each task is described in a following section. See the "IP Configuration Examples" section at the end of this chapter for broadcasting configuration examples.

## Enable Directed Broadcast-to-Physical Broadcast Translation

To enable forwarding of directed broadcasts on an interface where the broadcast becomes a physical broadcast, perform one of the tasks that follow. By default, this feature is enabled only for those protocols configured using the **ip forward-protocol** global configuration command. You can specify an access list to control which broadcasts are forwarded. When an access list is specified, only those IP packets permitted by the access list are eligible to be translated from directed broadcasts to physical broadcasts.

Perform either of these tasks in interface configuration mode as required for your network:

Task	Command
Enable directed broadcast-to-physical broadcast translation on an interface.	<b>ip directed-broadcast</b> [ <i>access-list-number</i> ]
Disable directed broadcast-to-physical broadcast translation on an interface.	<b>no ip directed-broadcast</b> [ <i>access-list-number</i> ]

## Forward UDP Broadcast Packets and Protocols

Network hosts occasionally employ User Datagram Protocol (UDP) broadcasts to determine address, configuration, and name information. If such a host is on a network segment that does not include a server, UDP broadcasts are normally not forwarded. You can remedy this situation by configuring the interface of your router to forward certain classes of broadcasts to a helper address. You can have more than one helper address per interface.

You can specify a UDP destination port to control which UDP services are forwarded. Multiple UDP protocols can be specified. You also can specify the Network Disk (ND) protocol, which is used by older diskless Sun workstations. By default, both UDP and ND forwarding are enabled if a helper

address has been defined for an interface. The description for the **ip forward-protocol** command in the *Router Products Command Reference* publication lists the services that are forwarded by default if you do not specify any ports.

To enable forwarding and to specify the destination address, perform the following task in interface configuration mode:

Task	Command
Enable forwarding and specify the destination address for forwarding UDP broadcast packets, including BOOTP.	<b>ip helper-address</b> <i>address</i>

To specify which protocols will be forwarded, perform the following task in global configuration mode:

Task	Command
Specify which protocols will be forwarded over which ports.	<b>ip forward-protocol</b> { <b>udp</b>   <b>nd</b> } [ <i>port</i> ]

See the “IP Configuration Examples” section at the end of this chapter for an example of how to configure helper addresses.

## Establish an Internet Broadcast Address

The router supports Internet broadcasts on both local- and wide-area networks. There are several ways to indicate an Internet broadcast address. Currently, the most popular way, and the default, is an address consisting of all ones (255.255.255.255), although the routers can be configured to generate any form of Internet broadcast address. Our routers also can receive and understand any form of Internet broadcast.

To set the router’s Internet broadcast address, perform the following task in interface configuration mode:

Task	Command
Establish a different broadcast address (other than 255.255.255.255).	<b>ip broadcast-address</b> [ <i>address</i> ]

If the router does not have nonvolatile memory, and you need to specify the broadcast address to use before the router has been configured, you have to change the Internet broadcast address by setting jumpers in the processor configuration register. Setting bit 10 causes the router to use all zeros. Bit 10 interacts with bit 14, which controls the network and subnet portions of the broadcast address. Setting bit 14 causes the router to include the network and subnet portions of its address in the broadcast address. Table 1-2 shows the combined effect of setting bits 10 and 14.

**Table 1-2 Configuration Register Settings for Broadcast Address Destination**

Bit 14	Bit 10	Address (<net><host>)
Out	Out	<ones><ones>
Out	In	<zeros><zeros>
In	In	<net><zeros>
In	Out	<net><ones>

Some router platforms allow the configuration register to be set through the software; see the “Loading System Images, Microcode Images, and Configuration Files” chapter for details. For other router platforms, the configuration register can only be changed through hardware; see the appropriate hardware installation and maintenance manual for your system.

## Flood IP Broadcasts

You can allow IP broadcasts to be flooded throughout your internetwork in a controlled fashion using the database created by the bridging spanning-tree protocol. Turning on this feature also prevents loops. The transparent bridging option must be in the routing software, and bridging must be configured on each interface that is to participate in the flooding, in order to support this capability. If an interface does not have bridging configured, it still will be able to receive broadcasts, but it will never forward broadcasts received on that interface, and it will never use that interface to send broadcasts received on a different interface.

The same packets that may be subject to the helper address you specify and forwarded to a single network can be flooded. Only one copy of the packet is put on each network segment.

Packets must meet the following criteria to be considered for flooding (these are the same conditions for IP helper addresses):

- The packet must be a MAC-level broadcast.
- The packet must be an IP-level broadcast.
- The packet must be a TFTP, DNS, Time, NetBIOS, ND, or BOOTP packet or a UDP protocol specified by the **ip forward-protocol udp** global configuration command.
- The packet’s time-to-live (TTL) value must be at least two.

A flooded UDP datagram is given the destination address you specified with the **ip broadcast-address** command on the output interface. The destination address can be set to any desired address. Thus, the destination address may change as the datagram propagates through the network. The source address is never changed. The TTL value is decremented.

After a decision has been made to send the datagram out on an interface (and the destination address possibly changed), the datagram is handed to the normal IP output routines and is therefore subject to access lists if they are present on the output interface.

To use the bridging spanning-tree database to flood UDP datagrams, perform the following task in global configuration mode:

Task	Command
Use the bridging spanning-tree database to flood UDP datagrams.	<b>ip forward-protocol spanning-tree</b>



If no actual bridging is desired, you can configure a type-code bridging filter that will deny all packet types from being bridged. Refer to the “Configuring Transparent Bridging” chapter of this manual for more information about using access lists to filter bridged traffic. The spanning-tree database is still available to the IP forwarding code to use for the flooding.

### Speed Up Flooding of UDP Datagrams

You can speed up flooding of UDP datagrams using the spanning-tree algorithm. Used in conjunction with the **ip forward-protocol spanning-tree** command, this feature boosts the performance of spanning tree-based UDP flooding by a factor of about four to five times. The feature, called *turbo flooding*, is supported over ARPA-encapsulated Ethernets, FDDI, and HDLC-encapsulated serials, but is not supported on Token Rings. As long as the Token Rings and the non-HDLC serials are not part of the bridge group being used for UDP flooding, turbo flooding will behave normally.

To enable turbo flooding, perform the following task in global configuration mode:

Task	Command
Use the bridging spanning-tree database to speed up flooding of UDP datagrams.	<b>ip forward-protocol turbo-flood</b>

## Configure IP Services

The Internet protocol suite offers a number of services that control and manage IP connections. Many of these services are provided by the Internet Control Message Protocol (ICMP). ICMP messages are sent by routers to hosts or other routers when a problem is discovered with the IP header. For detailed information on ICMP, see RFC 792.

This section summarizes the following tasks for configuring IP services:

- Disable ICMP protocol unreachable messages
- Disable ICMP redirect messages
- Understand path MTU discovery
- Set the MTU packet size
- Enable ICMP mask reply messages
- Disable IP source routing
- Configure simplex Ethernet interfaces

Each task is described in a following section.

See the “IP Configuration Examples” section at the end of this chapter for examples of ICMP services.

### Disable ICMP Protocol Unreachable Messages

If the router receives a nonbroadcast packet destined for itself that uses an unknown protocol, it sends an ICMP Protocol Unreachable message back to the source. Similarly, if the router receives a packet that it is unable to deliver to the ultimate destination because it knows of no route to the destination address, it sends an ICMP Host Unreachable message to the source. This feature is enabled by default.

You can disable this service by performing the following task in interface configuration mode:

Task	Command
Disable the sending of ICMP Protocol Unreachable and Host Unreachable messages.	<b>no ip unreachable</b>

## Disable ICMP Redirect Messages

Routers sometimes can become less than optimal. For example, it is possible for the router to be forced to resend a packet through the same interface on which it was received. If this happens, the router sends an ICMP *Redirect* message to the packet’s originator telling it that it is on a subnet directly connected to the router, and that it must forward the packet to another system on the same subnet. It does so because the originating host presumably could have sent that packet to the next hop without involving the router at all. The Redirect message instructs the sender to remove the router from the route and substitute a specified device representing a more direct path. This feature is enabled by default.

You can disable the sending of ICMP Redirect messages by performing the following task in interface configuration mode:


Task	Command
Disable the sending of ICMP Redirect messages to learn routes.	<b>no ip redirects</b>

## Understand Path MTU Discovery

Our routers support the IP Path MTU Discovery mechanism, as defined in RFC 1191. IP Path MTU Discovery allows a host to dynamically discover and cope with differences in the maximum allowable MTU size of the various links along the path. Sometimes a router is unable to forward a datagram because it requires fragmentation (the packet is larger than the MTU you set for the interface with the **ip mtu** command), but the “Don’t fragment” (DF) bit is set. The router sends a message to the sending host, alerting it to the problem. The host will have to fragment packets destined for the receiving interface so that they fit the smallest packet size of all the links along the path. This technique is shown in Figure 1-1.

Figure 1-1 Path MTU Discovery

Path MTU Discovery is useful when a link in a network goes down, forcing use of another, different MTU-sized link (and different routers). As shown in Figure 1-1, suppose one is trying to send IP packets over a network where the MTU in the first router is set to 1500 bytes, but then reaches a router where the MTU is set to 512 bytes. If the datagram’s “Don’t fragment” bit is set, the datagram would be dropped because the 512-byte router is unable to forward it. All packets larger than 512 bytes will be dropped in this case. The second router returns an ICMP Destination Unreachable



message to the source of the datagram with its Code field indicating “Fragmentation needed and DF set.” To support Path MTU Discovery, it also would include the MTU of the next-hop network link in the low-order bits of an unused header field.

Path MTU Discovery also is useful when a connection is first being established and the sender has no information at all about the intervening links. It is always advisable to use the largest MTU that the links will bear; the larger the MTU, the fewer packets the host needs to send.

---

**Note** MTU Discovery is a process initiated by end hosts. If an end host does not support MTU Discovery, a router will have no mechanism available to avoid fragmenting datagrams generated by the end host.

---

The Cisco 7000, AGS+, and Cisco 4000 routers support fast switching of IP packets between Ethernet and FDDI interfaces. When packets are being sent from FDDI to Ethernet interfaces and you are not using Path MTU Discovery, FDDI packets with data lengths larger than 1500 bytes will be fragmented into multiple Ethernet packets. This will slow performance. If the majority of your router traffic travels off the FDDI ring, you might want to either lower the MTU size on your host FDDI interfaces to 1500 bytes or run IP Path MTU Discovery on your hosts.

Because the CTR card does not support the switching of frames larger than 4472 bytes, some interoperability problems may occur if CTR cards are intermixed with other Token Ring cards on the same network. You can minimize this by setting lower (and the same) IP maximum packet sizes (MTUs) for all devices on the network with the **ip mtu** interface command.

## Set the MTU Packet Size

All interfaces have a default maximum transmission unit (MTU) packet size. You can adjust the IP MTU size so that if an IP packet exceeds the MTU set for a router’s interface, the router will fragment it.

Changing the MTU value (with the **mtu** interface configuration command) can affect the IP MTU value. If the current IP MTU value is the same as the MTU value, and you change the MTU value, the IP MTU value will be modified automatically to match the new MTU. However, the reverse is not true; changing the IP MTU value has no effect on the value for the **mtu** interface configuration command.

Also, all devices on a physical medium must have the same protocol MTU in order to operate.

To set the MTU packet size for a specified interface, perform the following task in interface configuration mode:

Task	Command
Set the IP MTU packet size for an interface.	<b>ip mtu bytes</b>

## Enable ICMP Mask Reply Messages

Occasionally, network devices need to know the subnet mask for a particular subnetwork in the internetwork. To achieve this information, such devices can send ICMP Mask Request messages. These messages are responded to by ICMP Mask Reply messages from devices that have the requested information. The router can respond to ICMP Mask Request messages if this function is enabled.

To enable the sending of ICMP Mask Reply messages, perform the following task in interface configuration mode:

Task	Command
Enable the sending of ICMP Mask Reply messages.	<b>ip mask-reply</b>

## Disable IP Source Routing

The router examines IP header options on every packet. It supports the IP header options *Strict Source Route*, *Loose Source Route*, *Record Route*, and *Time Stamp*, which are defined in RFC 791. If the router finds a packet with one of these options enabled, it performs the appropriate action. If it finds a packet with an invalid option, it sends an ICMP Parameter Problem message to the source of the packet and discards the packet.

IP provides a provision allowing the source IP host to specify a route through the IP network. This provision is known as *source routing*. Source routing is specified as an option in the IP header. If source routing is specified, the router forwards the packet according to the specified source route. This feature is employed when you want to force a packet to take a certain route through the network. The default is to perform source routing.

You can disable IP source-route header options by performing the following task in global configuration mode:

Task	Command
Cause the router to discard any IP datagram containing a source-route option.	<b>no ip source-route</b>

## Configure Simplex Ethernet Interfaces

You can configure simplex Ethernet interfaces. This feature is useful for setting up dynamic IP routing over a simplex circuit; that is, a circuit that receives only or transmits only. When a route is learned on a receive-only interface, the interface designated as the source of the route is converted

to the interface you specify. When packets are routed out this specified interface, they are sent to the IP address of the source of the routing update. To reach this IP address on a transmit-only Ethernet link, a static ARP entry mapping this IP address to the hardware address of the other end of the link is required.

To assign a transmit interface to a receive-only interface, perform the following task in interface configuration mode:

Task	Command
Assign a transmit interface to a receive-only interface.	<b>transmit-interface</b> <i>interface-name</i>

See the “IP Configuration Examples” section at the end of this chapter for an example of configuring a simplex Ethernet interface.

## Filter IP Packets

Packet filtering helps control packet movement through the network. Such control can help limit network traffic and restrict network use by certain users or devices. To permit or deny packets from crossing specified router interfaces, we provide *access lists*.

You can use access lists in several ways:

- To control the transmission of packets on an interface
- To control virtual terminal line access
- To restrict contents of routing updates

This section summarizes how to create access lists and how to apply them.

See the “IP Configuration Examples” section at the end of this chapter for examples of configuring access lists.

## Create an Access List

An access list is a sequential collection of permit and deny conditions that apply to IP addresses. The router tests addresses against the conditions in an access list one by one. The first match determines whether the router accepts or rejects the address. Because the router stops testing conditions after the first match, the order of the conditions is critical. If no conditions match, the router rejects the address.

The two steps involved in using access lists are as follows:

**Step 1** Create an access list by specifying an access list number and access conditions.

**Step 2** Apply the access list to interfaces.

These steps are described in the next sections.

## Create Standard and Extended Access Lists

The software supports two styles of access lists for IP:

- Standard IP access lists use source addresses for matching operations.

- Extended IP access lists use source and destination addresses for matching operations, as well as optional protocol type information for finer granularity of control.

To create either a standard or an extended access list, perform one of the following tasks in global configuration mode:

Task	Command
Define an IP access list number and the access conditions.	<b>access-list</b> <i>access-list-number</i> { <b>permit</b>   <b>deny</b> } <i>source</i> [ <i>source-mask</i> ]
Define an extended IP access list number and the access conditions.	<b>access-list</b> <i>access-list-number</i> { <b>permit</b>   <b>deny</b> } <i>protocol source source-mask destination destination-mask</i> [ <i>operator operand</i> ] [ <b>established</b> ]

After an access list is created initially, any subsequent additions (possibly entered from the terminal) are placed at the *end* of the list. In other words, you cannot selectively add or remove access list command lines from a specific access list.

---

**Note** Keep in mind when making the standard access list that by default, the end of the access list contains an implicit deny statement for *everything* if it did not find a match before reaching the end. Further, if you omit the mask from an associated IP host address access list specification, 0.0.0.0 is assumed to be the mask.

---

Refer to the “IP Configuration Examples” section at the end of this chapter for examples of implicit masks.

### Apply an Access List to an Interface

After an access list is created, you can apply it to one or more interfaces. Access lists can be applied on *either* outbound or inbound interfaces. The following two tables show how this task is accomplished for both terminal lines and network interfaces.

Perform the following task in line configuration mode:

Task	Command
Restrict incoming and outgoing connections between a particular virtual terminal line (into a device) and the addresses in an access list.	<b>access-class</b> <i>access-list-number</i> { <b>in</b>   <b>out</b> }

Perform the following task in interface configuration mode:

Task	Command
Control access to an interface.	<b>ip access-group</b> <i>access-list-number</i> { <b>in</b>   <b>out</b> }

---

**Note** You cannot use extended access lists when you have autonomous switching enabled.

---

For inbound access lists, after receiving a packet, the router checks the source address of the packet against the access list. If the access list permits the address, the router continues to process the packet. If the access list rejects the address, the router discards the packet and returns an ICMP *Host Unreachable* message.

For outbound access lists, after receiving and routing a packet to a controlled interface, the router checks the source address of the packet against the access list. If the access list permits the address, the router transmits the packet. If the access list rejects the address, the router discards the packet and returns an ICMP *Host Unreachable* message.

When you apply an access list (standard or extended) that has not yet been defined to an interface, the router will act as if the access list has not been applied to the interface and will accept all packets. Remember this behavior if you use undefined access lists as a means of security in your network.

---

**Note** Set identical restrictions on all the virtual terminal lines, because a user can attempt to connect to any of them.

---

## Configure IP Security Options

Our IP Security Option (IPSO) support addresses both the basic and extended security options as described in RFC 1108. Our implementation is only minimally compliant with RFC 1108, because our router only accepts and generates a four-byte IPSO. IPSO is generally used to comply with the U.S. Government's DoD security policy.

IPSO provides the following features:

- Defines security level on a per-interface basis
- Defines single-level or multilevel interfaces
- Provides a label for incoming packets
- Strips labels on a per-interface basis
- Reorders options to put any basic security options first
- Accepts or rejects messages with extended security options

Two general steps are required to configure IPSO-level security:

**Step 1** Enable IPSO and set the security classifications.

**Step 2** Select the security levels.

See the following sections for descriptions of each task. See the "IP Commands" chapter of the *Router Products Command Reference* publication for more specific details on the security options commands. See the "IP Configuration Examples" section at the end of this chapter for an example of configuring IPSO options.

## Enable IPSO and Set the Security Classifications

To enable IPSO and set security classifications on an interface, perform either of the following tasks in interface configuration mode:

Task	Command
Set an interface to the requested IPSO classification and authorities.	<b>ip security dedicated</b> <i>level authority [authority ...]</i>
or	
Set an interface to the requested IPSO range of classifications and authorities.	<b>ip security multilevel</b> <i>level1 [authority1 ...] to level2 authority2 [authority2 ...]</i>

Use the **no ip security** command to reset an interface to its default state.

## Select the Security Levels

To select the various IPSO security levels, perform any of the following optional tasks in interface configuration mode:

Task	Command
Enable an interface to ignore the authorities field of all incoming packets.	<b>ip security ignore-authorities</b>
Classify packets that have no IPSO with an implicit security label.	<b>ip security implicit-labelling</b> [ <i>level authority [authority ...]</i> ]
Accept packets on an interface that has an extended security option present.	<b>ip security extended-allowed</b>
Ensure that all packets leaving the router on an interface contain a basic security option.	<b>ip security add</b>
Remove any basic security option that may be present on a packet leaving the router through an interface.	<b>ip security strip</b>
Prioritize security options on a packet.	<b>ip security first</b>
Treat as valid any packets that have Reserved1 through Reserved4 security levels.	<b>ip security reserved-allowed</b>

## Default Values for Minor Keywords

In order to fully comply with IPSO, the default values for the minor keywords have become complex. Default value usages include the following:

- The default for all of the minor keywords is *off*, with the exception of **implicit-labelling** and **add**.
- The default value of **implicit-labelling** is *on* if the interface is unclassified Genser; otherwise it is *off*.
- The default value for **add** is *on* if the interface is not “unclassified Genser”; and otherwise it is *off*.



Table 1-3 provides a list of all default values.

**Table 1-3 Default Security Keyword Values**

Interface Type	Level	Authority	Implicit Labeling	Add IPSO
None	None	None	On	Off
Dedicated	Unclassified	Genser	On	Off
Dedicated	Any	Any	Off	On
Multilevel	Any	Any	Off	On

The default value for any interface is “dedicated, unclassified Genser.” Note that this implies implicit labeling. This may seem unusual, but it makes the system entirely transparent to packets without options. This is the setting generated when you specify the **no ip security** interface configuration command.

## Configure IP Accounting

Our IP accounting support provides basic IP accounting functions. By enabling IP accounting, users can see the number of bytes and packets switched through the router on a source and destination IP address basis. Only transit IP traffic is measured and only on an outbound basis; traffic generated by the router or terminating in the router is not included in the accounting statistics. To maintain accurate accounting totals, the router software maintains two accounting databases: an active and a checkpointed database.

You enable IP accounting on a per-interface basis. The following two tables list the tasks to configure IP accounting functions.

Perform the following task for each interface in interface configuration mode:

Task	Command
Enable IP accounting.	<b>ip accounting</b>

Perform the following tasks in global configuration mode:

Task	Command
Enable IP accounting for transit traffic outbound on an interface.	<b>ip accounting-threshold</b> <i>threshold</i>
Filter accounting information for hosts.	<b>ip accounting-list</b> <i>ip-address mask</i>
Control the number of transit records that will be stored in the IP accounting database.	<b>ip accounting-transits</b> <i>count</i>

Use the EXEC command **show ip accounting** to display the active accounting database. To display the checkpointed database, use the **show ip accounting checkpoint** EXEC command. The **clear ip accounting** EXEC command clears the active database and creates the checkpointed database.

## Configure Performance Parameters

You can tune IP performance in several ways using the following methods:

- Compress TCP packet headers
- Set the TCP connection attempt time
- Enable fast switching
- Enable autonomous switching
- Control route cache invalidation

### Compress TCP Packet Headers

You can compress the headers of your TCP/IP packets in order to reduce their size, thereby increasing performance. Header compression is particularly useful on networks with a large percentage of small packets, such as those supporting many Telnet connections. This feature only compresses the TCP header, so it has no effect on UDP packets or other protocol headers. The TCP header compression technique, described fully in RFC 1144, is supported on serial lines using HDLC or PPP encapsulation. You must enable compression on both ends of a serial connection.

You can optionally specify outgoing packets to be compressed only if TCP incoming packets on the same interface are compressed. If you do not specify this option (with the **passive** keyword), the router will compress all traffic. The default is no compression.

You also can specify the total number of header compression connections that can exist on an interface. You should configure one connection for each TCP connection through the specified interface.

To enable compression, perform either of the following optional tasks in interface configuration mode:

Task	Command
Enable TCP header compression.	<b>ip tcp header-compression</b> [ <b>passive</b> ]
Specify the total number of header compression connections that can exist on an interface.	<b>ip tcp compression-connections</b> <i>number</i>

---

**Note** When compression is enabled, fast switching is disabled. Fast processors can handle several fast interfaces, such as T1s, that are running header compression. However, you should think carefully about your network’s traffic characteristics before compressing TCP headers. You may want to use the monitoring commands to help compare network utilization before and after enabling header compression.

---

### Set the TCP Connection Attempt Time

You can set the amount of time the router will wait to attempt to establish a TCP connection. In previous versions of router software, the system would wait a fixed 30 seconds when attempting to do so. This amount of time is not sufficient in networks that have dialup asynchronous connections, such as a network consisting of Dial-on-Demand links that are implemented over modems, because it will affect your ability to Telnet over the link (from the router) if the link must be brought up.

Because this is a host parameter, it does not pertain to traffic going *through* the router, just to traffic originated *at* the router.

To set the TCP connection attempt time, perform the following task in global configuration mode:

Task	Command
Set the amount of time the router will wait to attempt to establish a TCP connection.	<b>ip tcp synwait-time</b> <i>seconds</i>

## Enable Fast Switching

Fast switching involves the use of a high-speed switching cache for IP routing. With fast switching, destination IP addresses are stored in the high-speed cache so that some time-consuming table lookups need not be done. Our routers generally offer better packet transfer performance when fast switching is enabled. However, if you have bursts of traffic, or if slow-speed serial links (64K and below) are being fed from higher-speed media such as T1 or Ethernet, then disabling fast switching can reduce the packet drop rate to some extent. Fast switching allows outgoing packets to be load balanced on a *per-destination* basis.

To enable or disable fast switching, perform the following tasks in interface configuration mode:


Task	Command
Enable fast-switching (use of a high-speed route cache for IP routing).	<b>ip route-cache</b>
Disable fast switching and enable load balancing on a per-packet basis.	<b>no ip route-cache</b>

## Enable Fast Switching on the Same Interface

You can enable IP fast switching when the input and output interfaces are the same interface. This normally is not recommended, though it is useful when you have partially meshed media, such as Frame Relay. You could use this feature on other interfaces, although it is not recommended because it would interfere with redirection.

Figure 1-2 illustrates a scenario where this is desirable. Router A has a DLCI to Router B, and Router B has a DLCI to Router C. There is no DLCI between Routers A and C; traffic between them must go in and out of Router B through the same interface.

Figure 1-2 IP Fast Switching on the Same Interface



Use the following IP interface configuration command to allow IP fast switching on the same interface:

Task	Command
Enable the fast switching of packets out of the same interface on which they arrived.	<b>ip route-cache same-interface</b>

### Enable IP Autonomous Switching

Autonomous switching contributes to faster packet processing by allowing the high-speed router bus to switch packets independently, without interrupting the system processor. This feature works only in Cisco 7000 or AGS+ systems with high-speed network controller cards, and with a switch processor or ciscoBus controller card running microcode Version 1.4 or later.

By default, IP autonomous switching is not enabled.

Perform any of the following tasks in interface configuration mode as needed for your network:

Task	Command
Enable both fast switching and autonomous switching.	<b>ip route-cache cbus</b>
Disable both fast switching and autonomous switching on an interface.	<b>no ip route-cache</b>
Disable only autonomous switching on an interface.	<b>no ip route-cache cbus</b>

**Note** Autonomous switching works only in high-end systems, such as the AGS+ and the Cisco 7000 series, with high-speed network controller cards, such as the CSC-HSCI, CSC-MEC, EIP, CSC-FCI, CSC-C2/FCIT, FIP, CSC-C2CTR, and TRIP, and with a ciscoBus controller card running microcode Version 1.4 or later. (See the information about microcode revisions in the release notes accompanying this publication for other microcode revision requirements.)

## Control Route Cache Invalidation

The high-speed route cache used by IP fast switching and autonomous switching is invalidated when the IP routing table changes. By default, the invalidation of the cache is delayed slightly to avoid excessive CPU load while the routing table is changing.

To control route cache invalidation, perform the following tasks in global configuration mode as needed for your network:

Task	Command
Allow immediate invalidation of the cache.	<b>no ip cache-invalidate-delay</b>
Delay invalidation of the cache.	<b>ip cache-invalidate-delay</b> [ <i>minimum maximum quiet threshold</i> ]

**Note** This task normally should not be necessary. It should be performed only under the guidance of technical staff. Incorrect configuration can seriously degrade the performance of your router.

## Configure IP over WANs

You can configure IP over X.25, SMDS, Frame Relay, and DDR networks. To do this, configure the appropriate address mappings, as described in the chapters “Configuring X.25 and LAPB,” “Configuring Frame Relay,” and “Configuring SMDS.”

## Monitor and Maintain the IP Network

You can perform the following tasks to monitor and maintain your network:

- Clear caches, tables, and databases
- Display system and network statistics

### Clear Caches, Tables, and Databases

You can remove all contents of a particular cache, table, or database. Clearing a cache, table, or database can become necessary when the contents of the particular structure have become or are suspected to be invalid.

The following table lists the tasks associated with clearing caches, tables, and databases. All are performed in EXEC mode.

Task	Command
Clear the IP ARP cache and the fast-switching cache.	<b>clear arp-cache</b>

Task	Command
Remove one or all entries from the host name and address cache.	<b>clear host</b> { <i>name</i>   *}
Clear the active IP accounting or checkpointed database when IP accounting is enabled.	<b>clear ip accounting</b> [ <b>checkpoint</b> ]
Remove one or more routes from the IP routing table.	<b>clear ip route</b> { <i>network</i> [ <i>mask</i> ]   *}

## Display System and Network Statistics

You can display specific router statistics such as the contents of IP routing tables, caches, and databases. Information provided can be used to determine resource utilization and solve network problems. You also can display information about node reachability and discover the routing path that your router's packets are taking through the network.

These tasks are summarized in the table that follows. See the “IP Commands” chapter in the *Router Products Command Reference* for details about the commands listed in these tasks. Perform the following tasks in EXEC mode:

Task	Command
Display the contents of all current access lists.	<b>show access-lists</b>
Display the entries in the ARP table for the router.	<b>show arp</b>
Display the default domain name, style of lookup service, the name server hosts, and the cached list of host names and addresses.	<b>show hosts</b>
Display the active IP accounting or checkpointed database.	<b>show ip accounting [checkpoint]</b>
Display Internet addresses mapped to TCP ports (aliases).	<b>show ip aliases</b>
Display the IP ARP cache.	<b>show ip arp</b>
Display the routing table cache used to fast switch IP traffic.	<b>show ip cache</b>
Display the usability status of interfaces.	<b>show ip interface [interface unit]</b>
Display the masks used for network addresses and the number of subnets using each mask.	<b>show ip masks address</b>
Display the address of a default gateway.	<b>show ip redirects</b>
Display the current state of the routing table.	<b>show ip route [address [mask]]   [protocol]</b>
Display the current state of the routing table in summary form.	<b>show ip route summary</b>
Show statistics on TCP header compression.	<b>show ip tcp header-compression</b>
Display IP protocol statistics.	<b>show ip traffic</b>
Test network node reachability (privileged).	<b>ping [protocol] {host   address}</b>
Test network node reachability using a simple ping facility (user).	<b>ping [protocol] {host   address}</b>
Trace packet routes through the network (privileged).	<b>trace [destination]</b>
Trace packet routes through the network (user).	<b>trace ip destination</b>

## IP Configuration Examples

This section shows complete configuration examples for the most common configuration situations. The configuration examples provided include the following topics:

- Serial interfaces configuration example (page 15-32)
- Example of creating a network from separated subnets (page 15-32)
- Dynamic lookup example (page 15-33)
- Example of establishing IP domains (page 15-33)
- Example of configuring HP hosts on a network segment (page 15-33)
- Helper addresses example (page 15-33)

- Broadcasting examples (page 15-34)
- Customizing ICMP services example (page 15-35)
- Simplex Ethernet interfaces example (page 15-36)
- Access list examples (page 15-36)
- IPSO configuration examples (page 15-38)
- Ping command example (page 15-40)

### Serial Interfaces Configuration Example

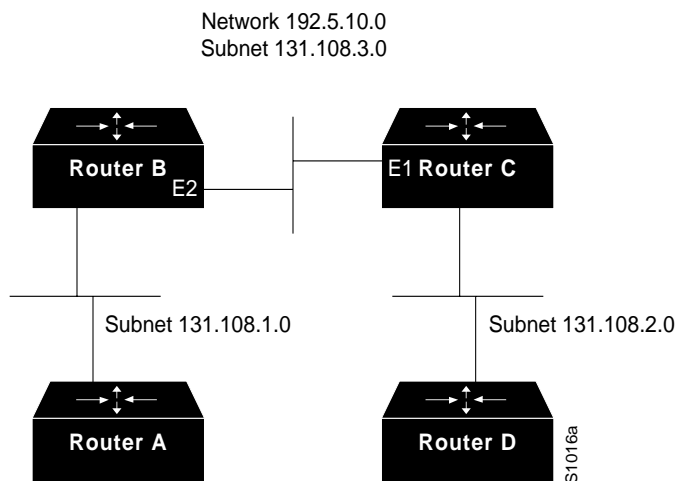
In the following example, the second serial interface (serial 1) is given Ethernet 0's address. The serial interface is unnumbered.

```
interface ethernet 0
ip address 145.22.4.67 255.255.255.0
interface serial 1
ip unnumbered ethernet 0
```

### Example of Creating a Network from Separated Subnets

In the following example, subnets 1 and 2 of network 131.108.0.0 are separated by a backbone, as shown in Figure 1-3. The two networks are brought into the same logical network through the use of secondary addresses.

Figure 1-3 Creating a Network from Separated Subnets



#### Configuration for Router B

```
interface ethernet 2
ip address 192.5.10.1 255.255.255.0
ip address 131.108.3.1 255.255.255.0 secondary
```



### Configuration for Router C

```
interface ethernet 1
ip address 192.5.10.2 255.255.255.0
ip address 131.108.3.2 255.255.255.0 secondary
```

## Dynamic Lookup Example

A cache of host name-to-address mappings is used by **connect**, **telnet**, **ping**, **trace**, **write net**, and **configure net EXEC** commands to speed the process of converting names to addresses. The commands used in this example specify the form of dynamic name lookup to be used. Static name lookup also can be configured.

The following example configures the host name-to-address mapping process for the router. IP DNS-based translation is specified, the addresses of the name servers are specified, and the default domain name is given.

```
! IP Domain Name System (DNS)-based host name-to-address translation is enabled
ip domain-lookup
! Specifies host 131.108.1.111 as the primary name server and host 131.108.1.2
! as the secondary server
ip name-server 131.108.1.111 131.108.1.2
! Defines cisco.com as the default domain name the router uses to complete
! unqualified host names
ip domain-name cisco.com
```

## Example of Establishing IP Domains

The example that follows establishes a domain list with several alternate domain names.

```
ip domain-list csi.com
ip domain-list telecomprog.edu
ip domain-list merit.edu
```

## Example of Configuring HP Hosts on a Network Segment

The following example has a network segment with Hewlett-Packard devices on it. The commands listed customize the router's first Ethernet port to respond to Probe name requests for bl4zip and to use Probe as well as ARP.


```
ip hp-host bl4zip 131.24.6.27
interface ethernet 0
arp probe
ip probe proxy
```

## Helper Addresses Example

In this example, one router is on network 191.24.1.0 and the other is on network 110.44.0.0, and you want to permit IP broadcasts from hosts on either network segment to reach both servers. Figure 1-4 illustrates how to configure the router that connects network 110 to network 191.24.1.

```
!
ip forward-protocol udp
!
interface ethernet 1
ip helper-address 110.44.23.7
interface ethernet 2
ip helper-address 191.24.1.19
```

Figure 1-4 IP Helper Addresses



S1017a


## Broadcasting Examples

Our routers support two types of broadcasting: directed broadcasting and flooding. A directed broadcast is a packet sent to a specific network or series of networks, while a flooded broadcast packet is sent to every network. The following examples show configurations for both types of broadcasting.

### Flooded Broadcast Example

Figure 1-5 shows a flooded broadcast packet being sent to every network. The packet that is incoming from interface E0 is flooded to interfaces E1, E2 and Serial 0.

Figure 1-5 IP Flooded Broadcast



S1009a

A directed broadcast address includes the network or subnet fields. For example, if the network address is 128.1.0.0, the address 128.1.255.255 indicates all hosts on network 128.1.0.0. This would be a directed broadcast. If network 128.1.0.0 has a subnet mask of 255.255.255.0 (the third octet is the subnet field), the address 128.1.5.255 specifies all hosts on subnet 5 of network 128.1.0.0—another directed broadcast.

### Flooding of IP Broadcasts Example

In the following example, flooding of IP broadcasts is enabled on all interfaces (two Ethernet and two serial). No bridging is permitted. The access list denies all protocols. No specific UDP protocols are listed by a separate **ip forward-protocol udp** interface configuration command, so the default protocols (TFTP, DNS, Time, NetBIOS, and BOOTP) will be flooded.

```
ip forward-protocol spanning-tree
bridge 1 protocol dec
access-list 201 deny 0x0000 0xFFFF
interface ethernet 0
bridge-group 1
bridge-group 1 input-type-list 201
interface ethernet 1
bridge-group 1
bridge-group 1 input-type-list 201
interface serial 0
bridge-group 1
bridge-group 1 input-type-list 201
interface serial 1
bridge-group 1
bridge-group 1 input-type-list 201
```

### Customizing ICMP Services Example


The example that follows changes some of the ICMP defaults for the first Ethernet interface. Disabling the sending of redirects could mean that you do not think your routers on this segment will ever have to send a redirect. Disabling the unreachable messages will have a secondary effect—it also will disable path MTU discovery, because path discovery works by having routers send unreachable messages. If you have a network segment with a small number of devices and an absolutely reliable traffic pattern—which could easily happen on a segment with a small number of little-used user devices—you would be disabling options that your router would be unlikely to use anyway.

```
interface ethernet 0
no ip unreachable
no ip redirects
```

## Simplex Ethernet Interfaces Example

The following is an example of configuring a simplex Ethernet interface. Figure 1-6 illustrates how to configure IP on two routers sharing transmit-only and receive-only Ethernet connections.

Figure 1-6 Simplex Ethernet Connections



### Configuration for Router 1

```
interface ethernet 0
ip address 128.9.1.1
!
interface ethernet 1
ip address 128.9.1.2
transmit-interface ethernet 0
!
!use show interfaces command to find router2-MAC-address-E0
!arp 128.9.1.4 router2-MAC-address-E0 arpa
```

### Configuration for Router 2

```
interface ethernet 0
ip address 128.9.1.3
transmit-interface ethernet 1
!
interface ethernet 1
ip address 128.9.1.4
!
!use show interfaces command to find router1-MAC-address-E1
!arp 128.9.1.1 router1-MAC-address-E1 arpa
```

## Access List Examples

In this example, network 36.0.0.0 is a Class A network whose second octet specifies a subnet; that is, its subnet mask is 255.255.0.0. The third and fourth octets of a network 36.0.0.0 address specify a particular host. Using access list 2, the router would accept one address on subnet 48 and reject all others on that subnet. The last line of the list shows that the router would accept addresses on all other network 36.0.0.0 subnets.

```
access-list 2 permit 36.48.0.3
access-list 2 deny 36.48.0.0 0.0.255.255
access-list 2 permit 36.0.0.0 0.255.255.255
interface ethernet 0
ip access-group 2
```

## Examples of Implicit Masks in Access Lists

IP access lists contain *implicit* masks. For instance, if you omit the mask from an associated IP host address access list specification, 0.0.0.0 is assumed to be the mask. Consider the following example configuration:

```
access-list 1 permit 0.0.0.0
access-list 1 permit 131.108.0.0
access-list 1 deny 0.0.0.0 255.255.255.255
```

For this example, the following masks are implied in the first two lines:

```
access-list 1 permit 0.0.0.0 0.0.0.0
access-list 1 permit 131.108.0.0 0.0.0.0
```

The last line in the configuration (using the deny keyword) can be left off, because IP access lists implicitly *deny* all other access. This is equivalent to finishing the access list with the following command statement:

```
access-list 1 deny 0.0.0.0 255.255.255.255
```

The following access list only allows access for those hosts on the three specified networks. It assumes that subnetting is not used; the masks apply to the host portions of the network addresses. Any hosts with a source address that does not match the access list statements will be rejected.

```
access-list 1 permit 192.5.34.0 0.0.0.255
access-list 1 permit 128.88.0.0 0.0.255.255
access-list 1 permit 36.0.0.0 0.255.255.255
! (Note: all other access implicitly denied)
```

To specify a large number of individual addresses more easily, you can omit the address mask that is all zeros from the **access-list** global configuration command. Thus, the following two configuration commands are identical in effect:

```
access-list 2 permit 36.48.0.3
access-list 2 permit 36.48.0.3 0.0.0.0
```

## Examples of Configuring Extended Access Lists

In this example, the first line permits any incoming TCP connections with destination ports greater than 1023. The second line permits incoming TCP connections to the SMTP port of host 128.88.1.2. The last line permits incoming ICMP messages for error feedback.

```
access-list 102 permit tcp 0.0.0.0 255.255.255.255 128.88.0.0 0.0.255.255 gt 1023
access-list 102 permit tcp 0.0.0.0 255.255.255.255 128.88.1.2 0.0.0.0 eq 25
access-list 102 permit icmp 0.0.0.0 255.255.255.255 128.88.0.0 255.255.255.255
interface ethernet 0
ip access-group 102
```

For another example of using an extended access list, suppose you have a network connected to the Internet, and you want any host on an Ethernet to be able to form TCP connections to any host on the Internet. However, you do not want Internet hosts to be able to form TCP connections to hosts on the Ethernet except to the mail (SMTP) port of a dedicated mail host.

SMTP uses TCP port 25 on one end of the connection and a random port number on the other end. The same two port numbers are used throughout the life of the connection. Mail packets coming in from the Internet will have a destination port of 25. Outbound packets will have the port numbers reversed. The fact that the secure system behind the router always will be accepting mail connections on port 25 is what makes it possible to separately control incoming and outgoing services. The access list can be configured on either the outbound or inbound interface.

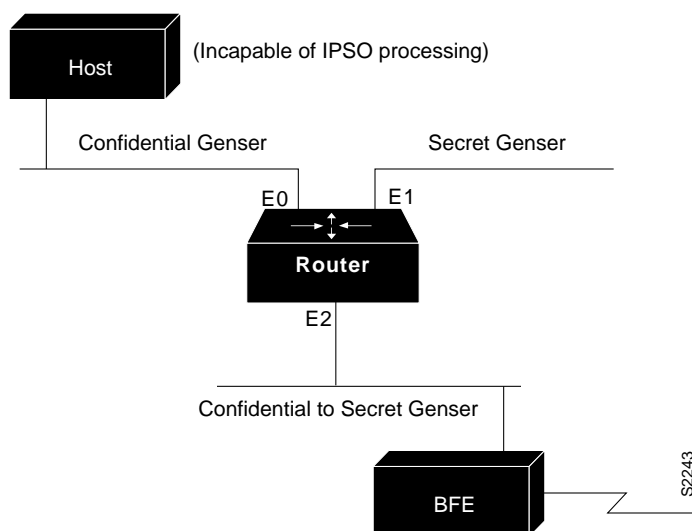
In the following example, the Ethernet network is a Class B network with the address 128.88.0.0, and the mail host's address is 128.88.1.2. The keyword **established** is used only for the TCP protocol to indicate an established connection. A match occurs if the TCP datagram has the ACK or RST bits set, which indicate that the packet belongs to an existing connection.

```
access-list 102 permit tcp 0.0.0.0 255.255.255.255 128.88.0.0 0.0.255.255 established
access-list 102 permit tcp 0.0.0.0 255.255.255.255 128.88.1.2 0.0.0.0 eq 25
interface ethernet 0
ip access-group 102
```

## IPSO Configuration Examples

In this example, three Ethernet interfaces are presented. These interfaces are running at security levels of Confidential Genser, Secret Genser, and Confidential to Secret Genser, as shown in Figure 1-7.

Figure 1-7 IPSO Security Levels



The following commands set up interfaces for the configuration in Figure 1-7:

```
interface ethernet 0
ip security dedicated confidential genser
interface ethernet 1
ip security dedicated secret genser
interface ethernet 2
ip security multilevel confidential genser to secret genser
```

It is possible for the setup to be much more complex.

In this next example, there are devices on Ethernet 0 that cannot generate a security option, and so must accept packets without a security option. These hosts do not understand security options; therefore, never place one on such interfaces. Furthermore, there are hosts on the other two networks that are using the extended security option to communicate information, so you must allow these to

pass through the system. Finally, there also is a host (a Blacker Front End) on Ethernet 2 that requires the security option to be the first option present, and this condition also must be specified. The new configuration follows.

```
interface ethernet 0
ip security dedicated confidential genser
ip security implicit-labelling
ip security strip
interface ethernet 1
ip security dedicated secret genser
ip security extended-allowed
!
interface ethernet 2
ip security multilevel confidential genser to secret genser
ip security extended-allowed
ip security first
```

## Ping Command Example

You can specify the router address to use as the source address for ping packets. In the following example it is 131.108.105.62.

```
Sandbox# ping
Protocol [ip]:
Target IP address: 131.108.1.111
Repeat count [5]:
Datagram size [100]:
Timeout in seconds [2]:
Extended commands [n]: yes
Source address: 131.108.105.62
Type of service [0]:
Set DF bit in IP header? [no]:
Data pattern [0xABCD]:
Loose, Strict, Record, Timestamp, Verbose[none]:
Type escape sequence to abort.
Sending 5, 100-byte ICMP Echos to 131.108.1.111, timeout is 2 seconds:
!!!!
Success rate is 100 percent, round-trip min/avg/max = 4/4/4 ms
```