

Configuring IP Services

This chapter describes how to configure optional IP services. For a complete description of the commands in this chapter, refer to the “IP Services Commands” chapter of the *Network Protocols Command Reference, Part 1*. To locate documentation of other commands that appear in this chapter, use the command reference master index or search online.

IP Services Task List

To configure optional IP services, complete any of the tasks in the following sections:

- Manage IP Connections
- Filter IP Packets
- Configure the Hot Standby Router Protocol
- Configure IP Accounting
- Configure Performance Parameters
- Configure IP over WANs
- Monitor and Maintain the IP Network

Remember that not all the tasks in these sections are required. The tasks you must perform will depend on your network and your needs.

At the end of this chapter, the examples in the “IP Services Configuration Examples” section illustrate how you might configure your network using IP.

Manage IP Connections

The IP suite offers a number of services that control and manage IP connections. ICMP provides many of these services. ICMP messages are sent by routers or access servers to hosts or other routers when a problem is discovered with the Internet header. For detailed information on ICMP, see RFC 792.

To manage various aspects of IP connections, perform the appropriate tasks in the following sections:

- Enable ICMP Protocol Unreachable Messages
- Enable ICMP Redirect Messages
- Enable ICMP Mask Reply Messages
- Understand Path MTU Discovery

- Set the MTU Packet Size
- Enable IP Source Routing
- Configure Simplex Ethernet Interfaces
- Configure a DRP Server Agent

See the “ICMP Services Example” section at the end of this chapter for examples of ICMP services.

Enable ICMP Protocol Unreachable Messages

If the Cisco IOS software 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 software 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 enable this service if it has been disabled by performing the following task in interface configuration mode:

Task	Command
Enable the sending of ICMP Protocol Unreachable and Host Unreachable messages.	ip unreachable

Enable ICMP Redirect Messages

Routes are sometimes 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 Cisco IOS software sends an ICMP Redirect message to the packet’s originator telling it that it is on a subnet directly connected to the receiving device, and that it must forward the packet to another system on the same subnet. The software does this because the originating host presumably could have sent that packet to the next hop without involving this device at all. The Redirect message instructs the sender to remove the receiving device from the route and substitute a specified device representing a more direct path. This feature is enabled by default. However, when Hot Standby Router Protocol is configured on an interface, ICMP Redirect messages are disabled by default for the interface.

You can enable the sending of ICMP Redirect messages if this feature was disabled by performing the following task in interface configuration mode:

Task	Command
Enable the sending of ICMP Redirect messages to learn routes.	ip redirects

Enable ICMP Mask Reply Messages

Occasionally, network devices must 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 Cisco IOS software 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.	<code>ip mask-reply</code>

Understand Path MTU Discovery

The Cisco IOS software supports 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 maximum transmission unit (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 Cisco IOS software sends a message to the sending host, alerting it to the problem. The host will have to fragment packets for the destination so that they fit the smallest packet size of all the links along the path. This technique is shown in Figure 14.

Figure 14 IP Path MTU Discovery



IP Path MTU Discovery is useful when a link in a network goes down, forcing the use of another, different MTU-sized link (and different routers). As shown in Figure 14, suppose a router is sending IP packets over a network where the MTU in the first router is set to 1,500 bytes, but the second router 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 are 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 IP Path MTU Discovery, it would also include the MTU of the next-hop network link in the low-order bits of an unused header field.

IP Path MTU Discovery is also 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 must send.

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

The Cisco 7000 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 IP 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 traffic travels off the FDDI ring, you might want to either lower the MTU size on your host FDDI interfaces to 1,500 bytes or run IP Path MTU Discovery on your hosts.

Because the CTR card does not support the switching of frames larger than 4,472 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 for all devices on the network with the **ip mtu** interface command.

To enable Path MTU Discovery for connections initiated by the router (when the router is acting as a host), see the section “Enable TCP Path MTU Discovery” later in this chapter.

Set the MTU Packet Size

All interfaces have a default MTU packet size. You can adjust the IP MTU size so that if an IP packet exceeds the MTU set for an interface, the Cisco IOS software 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.	ip mtu bytes

Enable IP Source Routing

The Cisco IOS software 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 software 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 that allows 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 software 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 enable IP source-route header options if they have been disabled by performing the following task in global configuration mode:

Task	Command
Enable IP source routing.	ip source-route

Configure Simplex Ethernet Interfaces

You can configure simplex Ethernet interfaces. This feature is useful for setting up dynamic IP routing over a simplex circuit (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.	transmit-interface <i>type number</i>

See the “Simplex Ethernet Interfaces Example” section at the end of this chapter for an example of configuring a simplex Ethernet interface.

Configure a DRP Server Agent

The Director Response Protocol (DRP) is a simple User Datagram Protocol (UDP)-based application developed by Cisco Systems. It enables Cisco’s DistributedDirector product to query routers (DRP Server Agents) in the field for Border Gateway Protocol (BGP) and Interior Gateway Protocol (IGP) routing table metrics between distributed servers and clients. DistributedDirector, a separate standalone product, uses DRP to transparently redirect end-user service requests to the topologically closest responsive server. DRP enables DistributedDirector to provide dynamic, scalable, and “network intelligent” Internet traffic load distribution between multiple geographically dispersed servers.

DRP Server Agents are border routers (or peers to border routers) that support the geographically distributed servers for which DistributedDirector service distribution is desired. Note that, because DistributedDirector makes decisions based on BGP and IGP information, all DRP Server Agents must have access to full BGP and IGP routing tables.

Refer to the *Cisco DistributedDirector 2501 Installation and Configuration Guide* or the *Cisco DistributedDirector 4700-M Installation and Configuration Guide* for information on how to configure DistributedDirector.

Perform the tasks in the following sections to configure and maintain the DRP Server Agent. The first task is required; the remaining tasks are optional.

- Enable the DRP Server Agent
- Limit the Source of DRP Queries
- Configure Authentication of DRP Queries and Responses

To monitor and maintain the DRP Server Agent, see the section “Monitor and Maintain the DRP Server Agent” later in this chapter.

For an example of configuring a DRP Server Agent, see the section “DRP Server Agent Example” at the end of this chapter.

Enable the DRP Server Agent

The DRP Server Agent is disabled by default. To enable it, perform the following task in global configuration mode:

Task	Command
Enable the DRP Server Agent.	ip drp server

Limit the Source of DRP Queries

As a security measure, you can limit the source of valid DRP queries. If a standard IP access list is applied to the interface, the Server Agent will respond only to DRP queries originating from an IP address in the list. If no access list is configured, the server agent will answer all queries.

If both an access group and a key chain (described in the next section) have been configured, both security mechanisms must allow access before a request is processed.

To limit the source of valid DRP queries, perform the following task in global configuration mode:

Task	Command
Control the sources of valid DRP queries by applying a standard IP access list.	ip drp access-group <i>access-list-number</i>

Configure Authentication of DRP Queries and Responses

Another available security measure is to configure the DRP Server Agent to authenticate DRP queries and responses. You define a key chain, identify the keys that belong to the key chain, and specify how long each key is valid. To do so, perform the following tasks beginning in global configuration mode:

Task	Command
Step 1 Identify which key chain to use to authenticate all DRP requests and responses.	ip drp authentication key-chain <i>name-of-chain</i>
Step 2 Identify a key chain (match the name configured in Step 1).	key chain <i>name-of-chain</i>
Step 3 In key chain configuration mode, identify the key number.	key <i>number</i>
Step 4 In key chain key configuration mode, identify the key string.	key-string <i>text</i>
Step 5 Optionally specify the time period during which the key can be received.	accept-lifetime <i>start-time</i> { infinite <i>end-time</i> duration <i>seconds</i> }
Step 6 Optionally specify the time period during which the key can be sent.	send-lifetime <i>start-time</i> { infinite <i>end-time</i> duration <i>seconds</i> }

When configuring your key chains and keys, keep these points in mind:

- The key chain configured for the DRP Server Agent in Step 1 must match the key chain in Step 2.
- The key configured in the primary agent in the remote router must match the key configured in the DRP Server Agent in order for responses to be processed.

- You can configure multiple keys with lifetimes, and the software will rotate through them. Note that the router needs to know the time. Refer to the NTP and calendar commands in the “Performing Basic System Management” chapter of the *Configuration Fundamentals Configuration Guide*.
- If authentication is enabled and multiple keys on the key chain happen to be active based on the **send-lifetime** values, the software uses only the first key it encounters for authentication.
- Use the **show key chain** command to display key chain information.

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 interfaces, we provide *access lists*.

You can use access lists in the following 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 IP access lists and how to apply them.

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

An access list is a sequential collection of permit and deny conditions that apply to IP addresses. The Cisco IOS software tests addresses against the conditions in an access list one by one. The first match determines whether the software accepts or rejects the address. Because the software stops testing conditions after the first match, the order of the conditions is critical. If no conditions match, the software 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 or name and access conditions.
- Step 2** Apply the access list to interfaces or terminal lines.

These steps are described in the next sections.

Create Standard and Extended Access Lists Using Numbers



Caution Release 11.1 and later releases introduced substantial changes to IP access lists. These extensions are backward compatible; migrating from a release earlier than Release 11.1 to the current image will convert your access lists automatically. However, previous releases are not upwardly compatible with these changes. Thus, if you save an access list with the current image and then use older software, the resulting access list will not be interpreted correctly. **This could cause you severe security problems.** Save your old configuration file before booting Release 11.1-or-later images.

The software supports the following 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, and optional protocol type information for finer granularity of control.
- Dynamic extended IP access lists grant access per user to a specific source or destination host basis through a user authentication process. In essence, you can allow user access through a firewall dynamically, without compromising security restrictions. Dynamic access lists and lock-and-key access are described in the “Configuring Traffic Filters” chapter of the *Security Configuration Guide*.

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

Task	Command
Define a standard IP access list using a source address and wildcard.	access-list <i>access-list-number</i> {deny permit} <i>source</i> [<i>source-wildcard</i>]
Define a standard IP access list using an abbreviation for the source and source mask of 0.0.0.0 255.255.255.255.	access-list <i>access-list-number</i> {deny permit} any

To create an extended access list using a number, perform one of the following tasks in global configuration mode:

Task	Command
Define an extended IP access list number and the access conditions. Use the log keyword to get access list logging messages, including violations. Use the log-input keyword to include input interface, source MAC address or VC in the logging output.	access-list <i>access-list-number</i> {deny permit} <i>protocol</i> <i>source</i> <i>source-wildcard</i> <i>destination</i> <i>destination-wildcard</i> [precedence <i>precedence</i>] [tos <i>tos</i>] [established] [log log-input]
Define an extended IP access list using an abbreviation for a source and source wildcard of 0.0.0.0 255.255.255.255, and an abbreviation for a destination and destination wildcard of 0.0.0.0 255.255.255.255.	access-list <i>access-list-number</i> {deny permit} <i>protocol</i> any any
Define an extended IP access list using an abbreviation for a source and source wildcard of <i>source</i> 0.0.0.0, and an abbreviation for a destination and destination wildcard of <i>destination</i> 0.0.0.0.	access-list <i>access-list-number</i> {deny permit} <i>protocol</i> host <i>source</i> host <i>destination</i>
Define a dynamic access list. For information about lock-and-key access, refer to the “Configuring Traffic Filters” chapter in the <i>Security Configuration Guide</i> .	access-list <i>access-list-number</i> [dynamic <i>dynamic-name</i> [timeout <i>minutes</i>]] {deny permit} <i>protocol</i> <i>source</i> <i>source-wildcard</i> <i>destination</i> <i>destination-wildcard</i> [precedence <i>precedence</i>] [tos <i>tos</i>] [established] [log log-input]

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 When creating an access list, remember 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, with standard access lists, if you omit the mask from an associated IP host address access list specification, 0.0.0.0 is assumed to be the mask.

Note Autonomous switching is not used when you have extended access lists.

Note The logging facility might drop some logging message packets if there are too many to be handled or if there is more than one logging message to be handled in 1 second. This behavior prevents the router from crashing due to too many logging packets. Therefore, the logging facility should not be used as a billing tool or an accurate source of the number of matches to an access list.

After creating an access list, you must apply it to a line or interface, as shown in the section “Apply the Access List to an Interface or Terminal Line” later in this chapter.

See the “Implicit Masks in Access Lists Examples” section at the end of this chapter for examples of implicit masks.

Create Standard and Extended Access Lists Using Names



Caution Named access lists will not be recognized by any software release prior to Cisco IOS Release 11.2.

You can identify IP access lists with an alphanumeric string (a name) rather than a number (1 to 199). This feature allows you to configure more than 99 standard IP and 100 extended IP access lists in a router. If you identify your access list with a name rather than a number, the mode and command syntax are slightly different. Currently, only packet and route filters can use a named list.

Consider the following before configuring named access lists:

- Access lists specified by name are not compatible with older releases.
- Not all access lists that accept a number will accept a name. Access lists for packet filters and route filters on interfaces can use a name.
- A standard access list and an extended access list cannot have the same name.
- Numbered access lists are also available, as described in the earlier section, “Create Standard and Extended Access Lists Using Numbers.”

To create a standard access list, perform the following tasks beginning in global configuration mode:

Task	Command
Step 1 Define a standard IP access list using a name.	ip access-list standard <i>name</i>

Task	Command
Step 2 In access-list configuration mode, specify one or more conditions allowed or denied. This determines whether the packet is passed or dropped.	deny { <i>source</i> [<i>source-wildcard</i>] any } or permit { <i>source</i> [<i>source-wildcard</i>] any }
Step 3 Exit access-list configuration mode.	exit

To create an extended access list using a name, perform the following tasks beginning in global configuration mode:

Task	Command
Step 1 Define an extended IP access list using a name.	ip access-list extended <i>name</i>
Step 2 In access-list configuration mode, specify the conditions allowed or denied. Use the log keyword to get access list logging messages, including violations. Use the log-input keyword to include input interface, source MAC address or VC in the logging output.	{ deny permit } <i>protocol source source-wildcard destination destination-wildcard</i> [precedence <i>precedence</i>] [tos <i>tos</i>] [established] [log log-input]
or	
Define an extended IP access list using an abbreviation for a source and source wildcard of 0.0.0.0 255.255.255.255, and an abbreviation for a destination and destination wildcard of 0.0.0.0 255.255.255.255.	{ deny permit } <i>protocol host source host destination</i>
or	
Define an extended IP access list using an abbreviation for a source and source wildcard of <i>source</i> 0.0.0.0, and an abbreviation for a destination and destination wildcard of <i>destination</i> 0.0.0.0.	dynamic <i>dynamic-name</i> [timeout <i>minutes</i>] { deny permit } <i>protocol source source-wildcard destination destination-wildcard</i> [precedence <i>precedence</i>] [tos <i>tos</i>] [established] [log log-input]
or	
Define a dynamic access list. For information about lock-and-key access, refer to the “Configuring Traffic Filters” chapter in the <i>Security Configuration Guide</i> .	

Note Autonomous switching is not used when you have extended access lists.

After you initially create an access list, you place any subsequent additions (possibly entered from the terminal) at the end of the list. In other words, you cannot selectively add access list command lines to a specific access list. However, you can use **no permit** and **no deny** commands to remove entries from a named access list.

Note When making the standard and extended access list, remember 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, with standard access lists, if you omit the mask from an associated IP host address access list specification, 0.0.0.0 is assumed to be the mask.

After creating an access list, you must apply it to a line or interface, as shown in the following section, “Apply the Access List to an Interface or Terminal Line.”

See the “Named Access List Example” section at the end of this chapter for an example of a named access list.

Apply the Access List to an Interface or Terminal Line

After you create an access list, 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 to accomplish this task for both terminal lines and network interfaces. Remember the following:

- When controlling access to a line, you must use a number.
- When controlling access to an interface, you can use a name or number.

Perform the following task in line configuration mode. Only numbered access lists can be applied to lines. Set identical restrictions on all the virtual terminal lines, because a user can attempt to connect to any of them.

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

Perform the following task in interface configuration mode:

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

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

When you apply an access list that has not yet been defined to an interface, the software 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.

Configure the Hot Standby Router Protocol

The Hot Standby Router Protocol provides high network availability because it routes IP traffic from hosts on Ethernet, FDDI, or Token Ring networks without relying on the availability of any single router.

This feature is useful for hosts that do not support a router discovery protocol (such as IRDP) and do not have the functionality to switch to a new router when their selected router reloads or loses power. Because existing TCP sessions can survive the *failover*, this protocol also provides a more transparent means of recovery for hosts that dynamically select a next hop for routing IP traffic.

When the Hot Standby Router Protocol is configured on a network segment, it provides a virtual MAC address and an IP address that is shared among routers in a group of routers that is running the Hot Standby Router Protocol. One of these devices is selected by the protocol to be the active router. The active router receives and routes packets destined for the group's MAC address. For n routers running the Hot Standby Router Protocol, there are $n + 1$ IP and MAC addresses assigned.

The Hot Standby Router Protocol detects when the designated active router fails, at which point a selected standby router assumes control of the Hot Standby group's MAC and IP addresses. A new standby router is also selected at that time.

Devices that are running the Hot Standby Router Protocol send and receive multicast UDP-based hello packets to detect router failure and to designate active and standby routers.

When the Hot Standby Router Protocol is configured on an interface, ICMP Redirect messages are disabled by default for the interface.

You can configure multiple Hot Standby groups on an interface, thereby making fuller use of the redundant routers. To do so, specify a group number for each Hot Standby command you configure for the interface.

Note Token Ring interfaces allow up to three Hot Standby groups each.

Note The Cisco 1000 series, Cisco 2500 series, Cisco 3000 series, and Cisco 4000 series that use Lance Ethernet hardware do not support multiple Hot Standby groups on a single Ethernet interface.

The Hot Standby Router Protocol is supported over Inter-Switch Link (ISL) encapsulation. Refer to the "Configuring Routing between VLANs with ISL Encapsulation" chapter in the *Cisco IOS Switching Services Configuration Guide*.

To enable the Hot Standby Router Protocol on an interface, perform the following task in interface configuration mode:

Task	Command
Enable the Hot Standby Router Protocol.	<code>standby [group-number] ip [ip-address] [secondary]</code>

To configure other Hot Standby group attributes that affect how the local router participates in the Hot Standby Router Protocol, perform one or more of the following tasks in interface configuration mode:

Task	Command
Configure the time between hello packets and the hold time before other routers declare the active router to be down.	standby [<i>group-number</i>] timers <i>hellotime holdtime</i>
Set the Hot Standby priority used in choosing the active router. The priority value range is from 1 to 255, where 1 denotes the lowest priority and 255 denotes the highest priority. Specify that, if the local router has priority over the current active router, the local router should attempt to take its place as the active router. Configure a preemption delay, after which the Hot Standby router preempts and becomes the active router.	standby [<i>group-number</i>] priority <i>priority</i> [preempt [delay <i>delay</i>]] or standby [<i>group-number</i>] [priority <i>priority</i>] preempt [delay <i>delay</i>]
Configure the interface to track other interfaces, so that if one of the other interfaces goes down, the device's Hot Standby priority is lowered.	standby [<i>group-number</i>] track <i>type number</i> [<i>interface-priority</i>]
Select an authentication string to be carried in all Hot Standby Router Protocol messages.	standby [<i>group-number</i>] authentication <i>string</i>
Configure Hot Standby Router Protocol to use the interface's burned-in address as its virtual MAC address instead of the preassigned MAC address (on Ethernet and FDDI) or the functional address (on Token Ring).	standby use-bia

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 Cisco IOS software on a source and destination IP address basis. Only transit IP traffic is measured and only on an outbound basis; traffic generated by the software or terminating in the software is not included in the accounting statistics. To maintain accurate accounting totals, the software maintains two accounting databases: an active and a checkpointed database.

Our IP accounting support also provides information identifying IP traffic that fails IP access lists. Identifying IP source addresses that violate IP access lists alerts you to possible attempts to breach security. The data also indicates that you should verify IP access list configurations. To make this feature available to users, you must enable IP accounting of access list violations using the **ip accounting access-violations** command. Users can then display the number of bytes and packets from a single source that attempted to breach security against the access list for the source destination pair. By default, IP accounting displays the number of packets that have passed access lists and were routed.

To enable IP accounting, perform one of the following tasks for each interface in interface configuration mode:

Task	Command
Enable basic IP accounting.	ip accounting

Task	Command
Enable IP accounting with the ability to identify IP traffic that fails IP access lists.	ip accounting access-violations

To configure other IP accounting functions, perform one or more of the following tasks in global configuration mode:

Task	Command
Set the maximum number of accounting entries to be created.	ip accounting-threshold <i>threshold</i>
Filter accounting information for hosts.	ip accounting-list <i>ip-address wildcard</i>
Control the number of transit records that will be stored in the IP accounting database.	ip accounting-transits <i>count</i>

To display IP access violations for a specific IP accounting database, perform the following task in EXEC mode:

Task	Command
Display IP access-violation information.	show ip accounting [checkpoint] access-violations

To display IP access violations, you must give the **access-violations** keyword on the command. If you do not specify the keyword, the command defaults to displaying the number of packets that have passed access lists and were routed. The access violations output displays the number of the access list failed by the last packet for the source and destination pair. The number of packets reveals how aggressive the attack is upon a specific destination.

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

To tune IP performance, complete any of the tasks in the following sections. To configure various switching options, refer to the “Switching Paths” chapter in the *Switching Services Configuration Guide*.

- Compress TCP Packet Headers
- Set the TCP Connection Attempt Time
- Enable TCP Path MTU Discovery
- Enable TCP Selective Acknowledgment
- Enable TCP Timestamp
- Set the TCP Maximum Read Size
- Set the TCP Window Size
- Set the TCP Outgoing Queue Size

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, the Cisco IOS software 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.	ip tcp header-compression [passive]
Specify the total number of header compression connections that can exist on an interface.	ip tcp compression-connections <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 might 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 Cisco IOS software will wait to attempt to establish a TCP connection. In previous versions of software, the system would wait a fixed 30 seconds when attempting to establish a connection. This amount of time is not sufficient in networks that have dial-up 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 the connection attempt time is a host parameter, it does not pertain to traffic going through the device, just to traffic originated at the device.

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

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

Enable TCP Path MTU Discovery

Path MTU Discovery is a method for maximizing the use of available bandwidth in the network between the end points of a TCP connection, and is described in RFC 1191. By default, this feature is disabled. Existing connections are not affected when this feature is turned on or off. To enable Path MTU Discovery, perform the following task in interface configuration mode:

Task	Command
Enable Path MTU Discovery.	ip tcp path-mtu-discovery [age-timer {minutes infinite}]

Customers using TCP connections to move bulk data between systems on distinct subnets would benefit most by enabling this feature. This might include customers using RSRB with TCP encapsulation, STUN, X.25 Remote Switching (also known as XOT or X.25 over TCP), and some protocol translation configurations.

The **ip tcp path-mtu-discovery** command is to enable Path MTU Discovery for connections initiated by the router when it is acting as a host. For a discussion of how the Cisco IOS software supports Path MTU Discovery when the device is acting as a router, see the section “Understand Path MTU Discovery” earlier in this chapter.

The age-timer is a time interval for how often TCP should re-estimate the Path MTU with a larger maximum segment size (MSS). The default path MTU Discovery age-timer is 10 minutes; its maximum is 30 minutes. You can turn off the age-timer by setting it to infinite.

Enable TCP Selective Acknowledgment

The TCP selective acknowledgment feature improves performance in the event that multiple packets are lost from one TCP window of data.

Prior to this feature, with the limited information available from cumulative acknowledgments, a TCP sender could learn about only one lost packet per round trip time. An aggressive sender could choose to retransmit packets early, but such retransmitted segments might have already been successfully received.

The TCP selective acknowledgment mechanism helps improve performance. The receiving TCP host returns selective acknowledgment packets to the sender, informing the sender of data that has been received. In other words, the receiver can acknowledge packets received out of order. The sender can then retransmit only the missing data segments (instead of everything since the first missing packet).

Prior to selective acknowledgment, if TCP lost packets 4 and 7 out of an 8-packet window, TCP would receive acknowledgment of only packets 1, 2, and 3. Packets 4 through 8 would have to be resent. With selective acknowledgment, TCP receives acknowledgment of packets 1, 2, 3, 5, 6, and 8. Only packets 4 and 7 have to be resent.

Refer to RFC 2018 for more detailed information on TCP selective acknowledgment.

The feature is used only when multiple packets are dropped within one TCP window. There is no performance impact when the feature is enabled but not used. To enable TCP selective acknowledgment, perform the following task in global configuration mode.

Task	Command
Enable TCP selective acknowledgment.	ip tcp selective-ack

Enable TCP Timestamp

The TCP timestamp option provides better TCP round-trip time measurements. Because the timestamps are always sent and echoed in both directions and the timestamp value in the header is always changing, TCP header compression will not compress the outgoing packet. To allow TCP header compression over a serial link, the TCP timestamp option is disabled.

Refer to RFC 1323 for more detailed information on TCP timestamp.

To enable TCP timestamp, perform the following task in global configuration mode.

Task	Command
Enable TCP timestamp.	ip tcp timestamp

If you want to use TCP header compression over a serial line, TCP timestamp and TCP selective acknowledgment must be disabled. Both features are disabled by default. To disable TCP selective acknowledgment once it is enabled, refer to the “Enable TCP Selective Acknowledgment” section earlier in this chapter.

Set the TCP Maximum Read Size

By default, for Telnet and rlogin, the maximum number of characters that TCP reads from the input queue at once is a very large number (the largest possible 32-bit positive number). We do not recommend that you change this value. However, you could change that value by performing the following task in global configuration mode.

Task	Command
Set the TCP maximum read size for Telnet or rlogin.	ip tcp chunk-size <i>characters</i>

Set the TCP Window Size

The default TCP window size is 2144 bytes. We recommend you keep the default value, unless you know your router is sending large packets (greater than 536 bytes). To change the default window size, perform the following task in global configuration mode.

Task	Command
Set the TCP window size.	ip tcp window-size <i>bytes</i>

Set the TCP Outgoing Queue Size

The default TCP outgoing queue size per connection is 5 segments if the connection has a TTY associated with it (like a Telnet connection). If there is no TTY connection associated with it, the default queue size is 20 segments. To change the 5-segment default value, perform the following task in global configuration mode.

Task	Command
Set the TCP outgoing queue size.	ip tcp queuemax <i>packets</i>

Configure IP over WANs

You can configure IP over X.25, SMDS, Frame Relay, and DDR networks. To do this, configure the address mappings, as described in the appropriate chapters of the *Wide-Area Networking Configuration Guide*.

Monitor and Maintain the IP Network

To monitor and maintain your network, perform the tasks in the following sections:

- Clear Caches, Tables, and Databases
- Monitor and Maintain the DRP Server Agent
- Clear the Access List Counters
- 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. Perform the following tasks as needed in EXEC mode:

Task	Command
Clear the active IP accounting or checkpointed database when IP accounting is enabled.	clear ip accounting [checkpoint]
Clear TCP statistics.	clear tcp statistics

Monitor and Maintain the DRP Server Agent

To monitor and maintain the DRP Server Agent, perform the following tasks in EXEC mode:

Task	Command
Clear statistics being collected on DRP requests and responses.	clear ip drp
Display information about the DRP Server Agent.	show ip drp

Clear the Access List Counters

The system counts how many packets pass each line of an access list; the counters are displayed by the **show access-lists** command. You can clear the counters of an access list by performing the following task in EXEC mode:

Task	Command
Clear the access list counters.	clear access-list counters <i>access-list-number</i> <i>name</i>

Display System and Network Statistics

You can display specific statistics such as the contents of IP routing tables, caches, and databases. The resulting information can be used to determine resource utilization and to solve network problems.

These tasks are summarized in the table that follows. See the “IP Services Commands” chapter in the *Network Protocols Command Reference, Part 1* for details about the commands listed in these tasks. Perform any of the following tasks in privileged EXEC mode:

Task	Command
Display the contents of one or all current access lists.	show access-lists [<i>access-list-number</i> <i>name</i>]
Display the contents of current IP access lists.	show ip access-list [<i>access-list-number</i> <i>name</i>]
Display the active IP accounting or checkpointed database.	show ip accounting [checkpoint]
Show statistics on TCP header compression.	show ip tcp header-compression
Display IP protocol statistics.	show ip traffic
Display the status of the standby router.	show standby [<i>interface</i> [<i>group</i>]] [brief]
Display TCP statistics.	show tcp statistics

IP Services Configuration Examples

The following sections provide IP configuration examples:

- ICMP Services Example
- Simplex Ethernet Interfaces Example
- DRP Server Agent Example
- Numbered Access List Examples
- Named Access List Example

ICMP Services Example


The example that follows changes some of the ICMP defaults for the first Ethernet interface 0. Disabling the sending of redirects could mean that you do not think your devices on this segment will ever have to send a redirect. Disabling the Unreachables messages will have a secondary effect—it also will disable IP Path MTU Discovery, because path discovery works by having the Cisco IOS software send Unreachables 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 device 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 15 illustrates how to configure IP on two routers sharing transmit-only and receive-only Ethernet connections.

Figure 15 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
```

DRP Server Agent Example

The following example enables the DRP Server Agent. Sources of DRP queries are limited by access list 1, which permits only queries from the host at 33.45.12.4. Authentication is also configured for the DRP queries and responses.

```
ip drp server
access-list 1 permit 33.45.12.4
ip drp access-group 1
ip drp authentication key-chain mktg
key chain mktg
 key 1
  key-string internal
exit
exit
```

Numbered Access List Examples

In the following 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 Cisco IOS software would accept one address on subnet 48 and reject all others on that subnet. The last line of the list shows that the software 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 in
```

Implicit Masks in Access Lists Examples

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
```

Extended Access List Examples

In the following 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 in
```

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 IP 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 in
```

Named Access List Example

The following configuration creates a standard access list named `Internet_filter` and an extended access list named `marketing_group`:

```
interface Ethernet0/5
 ip address 2.0.5.1 255.255.255.0
 ip access-group Internet_filter out
 ip access-group marketing_group in
...
ip access-list standard Internet_filter
 permit 1.2.3.4
 deny any
ip access-list extended marketing_group
 permit tcp any 171.69.0.0 0.0.255.255 eq telnet
 deny tcp any any
 permit icmp any any
 deny udp any 171.69.0.0 0.0.255.255 lt 1024
 deny ip any any log
```