Troubleshooting TCP/IP

This chapter presents protocol-related troubleshooting information for Transmission Control Protocol/Internet Protocol (TCP/IP) connectivity and performance problems.

The sections in this chapter focus on general TCP/IP problems and on routing problems related to the Routing Information Protocol (RIP), the Interior Gateway Routing Protocol (IGRP), Enhanced IGRP, Open Shortest Path First (OSPF), the Border Gateway Protocol (BGP), and the Hot Standby Router Protocol (HSRP). Each section describes a specific symptom, the problems that are likely to cause each symptom, and the solutions to those problems.

- TCP/IP: Local Host Cannot Access Remote Host
- TCP/IP: Routes Learned from Wrong Interface or Protocol
- TCP/IP: Routing Not Functioning Properly on New Interface
- TCP/IP: Host Connections Fail Using Certain Applications
- TCP/IP: Problems Forwarding BOOTP and Other UDP Broadcasts
- TCP/IP: Poor Performance
- RIP/IGRP: Routes Missing from Routing Table
- OSPF: Routers Not Establishing Neighbors
- OSPF: Routes Missing from Routing Table
- IP Enhanced IGRP: Routers Not Establishing Neighbors
- IP Enhanced IGRP: Routes Missing from Routing Table
- IP Enhanced IGRP: Router Stuck in Active Mode
- BGP: Routes Missing from Routing Table
- BGP: Routers Not Advertising Routes
- HSRP: Hosts Cannot Reach Remote Networks

The symptoms described in the following sections are generic in nature and pertain to general TCP/IP internetwork problems. However, when host configuration problems are discussed, they are addressed assuming UNIX end systems. Similar types of actions might be applicable for non-UNIX hosts, but the discussion does not specifically address non-UNIX end-station problems.

TCP/IP: Local Host Cannot Access Remote Host

Symptom: Hosts on one network cannot communicate with hosts on a remote network. The networks are separated by one or more routers and might include WAN or other links. There is one or more routing protocol running on the routers.

Table 5-1 outlines the problems that might cause this symptom and describes solutions to those problems.

Table 5-1 TCP/IP: Local Host Cannot Access Remote Host

Possible Problem	Solution
Default gateway is not specified or is misconfigured on local or remote host	If hosts are not running *routed*, a default gateway should be configured. **Step 1** Determine whether the local and remote hosts have a default gateway specification. Use the following UNIX command: `unix-host% netstat -rn` Check the output of this command for a default gateway specification. **Step 2** If the default gateway specification is incorrect, or if it is not present at all, you can change or add a default gateway using the following UNIX command at the local host: `unix-host% route add default address 1` where *address* is the IP address of the default gateway (the router local to the host). The value 1 indicates that the specified gateway is one hop away. You might need to reboot the host for this change to take effect. **Step 3** It is recommended that you specify a default gateway as part of the boot process. Specify the IP address of the gateway in the */etc/defaultrouter* UNIX host file. This filename might be different on your UNIX system. If you are working with a PC or a Macintosh, consult the corresponding documentation to determine how to set the default gateway.
Misconfigured or missing routed default routes	**Step 1** If the host is running routed, use the **netstat -rn** UNIX command to view the host's routing table. The entry with Destination "default" denotes the default route. **Step 2** The default route entry should point to the router which has the route to the remote host. If there is no default route entry, use the **route** UNIX command to manually configure the default gateway.
DNS[1] host table is incomplete	If the DNS host table is incomplete, the DNS cannot reply to some lookup requests. If the DNS receives a lookup request for a hostname that is not in its cache, it cannot reply to the request, and the client cannot establish a connection. **Step 1** At the UNIX prompt, enter the following command: `unix-host% host address` where *address* is the IP address of a server, router, or other network node. **Step 2** If the result of this command is "Host not found," but you can open the connection using the host's IP address rather than its name, try connecting to other hosts using their names. If connections to other hosts can be opened using their names, then the host table might be incomplete. Add hostname-to-address mappings to the DNS cache for every host on the network. **Step 3** If you cannot open any connections using host names, the DNS might not be up and running. For troubleshooting information, see the following problem, "DNS is not up and running."

Possible Problem	Solution
DNS is not up and running	If issuing the **host** command at the UNIX prompt returns a "Host not found" message, but you are able to open a connection using the host's IP address, the DNS might not be up and running. Consult the DNS software documentation or your system administrator for information on configuring and enabling the DNS.
Routing is not enabled on one or more routers	**Step 1** Use the **traceroute** EXEC command to isolate the problem router (or routers). **Step 2** When you find a suspect router, determine if routing is enabled on the router. Enter the **show ip route** privileged EXEC command to view the routing table. Examine the output to see if the routing table is populated with routing information. **Step 3** If routing information is not being exchanged (that is, if the output of the **show ip route** command shows no entries that were learned from a routing protocol), use the **show running-config** privileged EXEC command on the router. **Step 4** Look for a **router** global configuration command for the routing protocol that should be enabled. For example, if the router should be running IGRP, look for an entry such as the following: `router igrp 109` ` network 192.168.52.0` ` network 192.168.48.0` **Step 5** If routing is not enabled on the router (or routers), enable the proper routing protocol using the **router** global configuration command. **Step 6** In router configuration mode, enter the appropriate **network** commands to associate networks with the routing process, as applicable. For example, to enable IGRP routing for networks 193.166.66.0 and 193.168.25.0, enter the following configuration commands: `Router(config)# router igrp 109` `Router(config-router)# network 193.166.66.0` `Router(config-router)# network 193.168.25.0` For complete information on configuring specific IP routing protocols, see the Cisco IOS *Network Protocols Configuration Guide, Part 1* and *Network Protocols Command Reference, Part 1*.
Routing is misconfigured on one or more routers	Narrow the specific symptoms down and troubleshoot the problem using the procedures outlined later in this chapter. For example, check the routing tables on various routers using the **show ip route** privileged EXEC command. If you are running IGRP and there are routes missing from the routing table (that is, you see no routes to certain networks that you know are connected), refer to the section "RIP/IGRP: Routes Missing from Routing Table" later in this chapter.

1. DNS=Domain Name Service

TCP/IP: Routes Learned from Wrong Interface or Protocol

Symptom: Routes in the routing table were learned from the wrong interface or protocol. For example, networks that should be reached through one interface are shown in the routing table to be reachable through another interface instead. This problem occurs only in a multiprotocol environment (see the section "Split Horizon Example," later in this chapter).

Table 5-2 outlines the problems that might cause this symptom and describes solutions to those problems.

Table 5-2 TCP/IP: Routes Learned from Wrong Interface or Protocol

Possible Problem	Solution
Split horizon is disabled	**Step 1** Use the **show ip interface** privileged EXEC command on the remote router to see the router configuration. **Step 2** Make sure that split horizon is enabled. Check the output of the **show ip interface** command for the following: `Split horizon is enabled` **Step 3** If split horizon is not enabled, enter the **ip split-horizon** interface configuration command on the remote router interface. For example, to enable split horizon on serial interface 0, enter the following commands: `C4500(config)#interface s0` `C4500(config-if)#ip split-horizon` **Note:** The default split-horizon setting for all LAN interfaces is *enabled*. However, for WAN multipoint interfaces configured with X.25, Frame Relay, or SMDS encapsulation, the default split-horizon setting is *disabled*.

Split Horizon Example

Sometimes in a multipoint WAN environment it is desirable to leave split horizon disabled. However, steps should be taken to prevent routing information from being learned from the wrong interface or protocol. For example, in the environment shown in Figure 5-1, Router 2 might incorrectly receive information about RIP networks from Router 3 if the routers are not configured correctly.

Figure 5-1 Split Horizon Sample Network

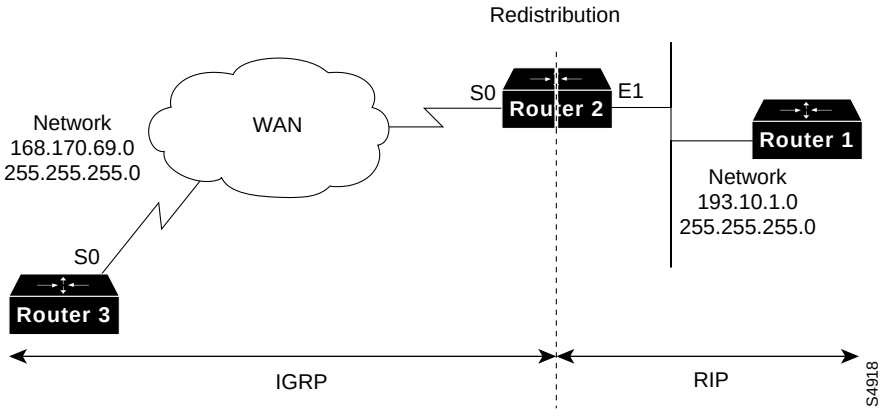

RIP routing information learned by Router 2 from Router 1 is redistributed into the IGRP domain. IGRP routing updates are sent to Router 3 from Router 2. If split horizon is disabled on Router 3, Router 3's updates to Router 2 will include information about network 193.10.1.0 (which was originally learned from RIP updates sent from Router 1 to Router 2).

Because IGRP routes by default are given a lower (better) administrative distance than RIP routes, Router 2 will route traffic to network 193.10.1.0 out serial interface 0 (towards Router 3) rather than out Ethernet interface 1 (towards Router 1).

Enabling split horizon on Router 3's serial interface prevents the router from advertising any of the RIP routes it has learned. However, in some cases, enabling split horizon is not desirable (for example, in a hub-and-spoke environment). In such a situation, route filtering using an input distribution list can be configured on Router 2's serial interface 0, as the following example shows:

```
Router_2(config)#router igrp 100
Router_2(config-router)#distribute-list 5 in
Router_2(config)#access-list 5 deny 193.10.1.0 255.255.255.0
Router_2(config)#access-list 5 permit 168.170.69.0 255.255.255.0
```

This distribution list specifically denies routing updates from Router 3 that advertise network 193.10.1.0, thus preventing Router 2 from learning information about this network from the wrong protocol and the wrong interface. Be sure to configure explicit **permit** statements for any traffic that you do want Router 2 to accept.

TCP/IP: Routing Not Functioning Properly on New Interface

Symptom: A new interface is added to a router, but when routing is configured it does not function properly on the new interface.

Table 5-3 outlines the problems that might cause this symptom and describes solutions to those problems.

Table 5-3 TCP/IP: Routing Not Functioning Properly on New Interface

Possible Problem	Solution
Interface or LAN protocol is down	**Step 1** Use the **show interfaces** privileged EXEC command to see whether the interface is "administratively down." **Step 2** If the interface is "administratively down," bring the interface up using the **no shutdown** interface configuration command. **Step 3** Use the **show interfaces** command again to see whether the interface is now up. **Step 4** If the interface is still down, there might be a hardware or media problem. See the procedures outlined in the "Troubleshooting Hardware and Booting Problems" and "Troubleshooting LAN Media Problems" chapters.
Misconfigured or missing **network** router configuration command	**Step 1** Use the **show running-config** privileged EXEC command to view the router configuration. **Step 2** Make sure that there is a **network** router configuration command specified for the network to which the interface belongs. For example, if you assign the new interface IP address 192.168.52.42, enter the following commands to enable RIP on the interface: `c4500(config)#router rip` `c4500(config-router)#network 192.168.52.0` Make sure that process IDs, addresses, and other variables are properly specified for the routing protocol you are using. For more information, refer to the Cisco IOS configuration guides and command references.
No active interfaces are configured with an IP address (OSPF only)	OSPF uses an IP address on the router as its router ID. Therefore, to configure the OSPF protocol on a router, you need at least one active interface configured with an IP address. If there is no active interface with an IP address, the router will return the following error: `2509(config)#router ospf 100` `2509(config)#` `OSPF: Could not allocate router id` **Step 1** Use the **show ip interfaces** privileged EXEC command on the router to make sure there is a router interface that is up and configured with an IP address. **Step 2** If there is no active interface with an IP address, configure an interface with the **ip address** interface configuration command. If necessary, use the **no shutdown** interface configuration command to bring an interface up.

TCP/IP: Host Connections Fail Using Certain Applications

Symptom: Connection attempts using some applications are successful, but attempts using other applications fail. For instance, you might be able to **ping** a host successfully, but Telnet connections fail.

Table 5-4 outlines the problems that might cause this symptom and describes solutions to those problems.

Table 5-4 TCP/IP: Host Connections Fail Using Certain Protocols

Possible Problem	Solution
Misconfigured access lists or other filters	**Step 1** Use the **show running-config** command to check each router in the path. See if there are IP access lists configured on the router. **Step 2** If there are IP access lists enabled on the router, disable them using the appropriate commands. For example, to disable input access list 80, enter the following command: `C4000(config-if)#no ip access-group 80 in` **Step 3** After disabling all of the access lists on the router, determine if the application in question operates normally. **Step 4** If the application operates normally, an access list is probably blocking traffic. **Step 5** To isolate the problem list, enable access lists one at a time until the application no longer functions. Check the problem access list to see if it is filtering traffic from any TCP or UDP ports. **Step 6** If the access list denies specific TCP or UDP ports, make sure that it does not deny the port used by the application in question (such as TCP port 23 for Telnet). Enter explicit **permit** statements for those ports used by applications you want to have functional. **Step 7** If you altered an access list, enable the list to see if the application can still operate normally. **Step 8** If the application operates normally, perform the preceding steps to isolate any other problem access lists until the application operates correctly with all access lists enabled. For more information about misconfigured access lists, see the section "Misconfigured Access List Example" later in this chapter. For more information on configuring access lists, see the Cisco IOS configuration guides and command references.

Misconfigured Access List Example

Misconfigured access lists can cause connectivity and performance problems. In the environment shown in Figure 5-2, the network administrator can successfully reach Router Z from Router X using the **telnet** and **ping** commands. However, when attempts are made to trace the route using the **trace** command, the connection fails.

Figure 5-2 Misconfigured Access List Sample Network

When examining the configuration of Router Y, the network administrator finds the following extended access list configured on the router:

```
C4500#show ip access-lists
Extended IP access list 101
    permit tcp any any eq telnet
    permit icmp any any
C4500#show running-config
[...]
interface Serial0
 ip address 192.168.54.92 255.255.255.0
 ip access-group 101 out
[...]
```

The access list permits only ICMP (used by the **ping** application) and TCP (used by the Telnet application) traffic to pass serial interface 0. Any traffic destined for UDP ports, including the default ports used by the trace application (UDP ports 33434 and above), is implicitly denied.

To allow trace traffic to pass through Router Y, the network administrator makes the following change to the access list:

```
C4500#configure terminal
C4500(config)#access-list 101 permit udp any any gt 33433
C4500(config)#^Z
C4500#
%SYS-5-CONFIG_I: Configured from console by console
C4500#show ip access-lists
Extended IP access list 101
    permit tcp any any eq telnet
    permit icmp any any
    permit udp any any gt 33433
C4500#
```

TCP/IP: Problems Forwarding BOOTP and Other UDP Broadcasts

Symptom: Problems occur when forwarding BOOTP or other UDP broadcast packets. UDP broadcasts sent from network hosts are not forwarded by routers. Diskless workstations cannot boot.

Table 5-5 outlines the problems that might cause this symptom and describes solutions to those problems.

Table 5-5 TCP/IP: Problems Forwarding BOOTP and Other UDP Broadcasts

Possible Problem	Solution
Missing or misconfigured **ip helper-address** specification	**Step 1** Use the **debug ip udp** privileged EXEC command on the router that should be receiving packets from the host. Check the output of the command to see if packets are being received from the host. **Caution:** This **debug** command can use considerable CPU cycles on the router. Do not enable it if your network is heavily congested. You can attach a protocol analyzer to see if UDP broadcasts are being received from the host if your network is congested. **Step 2** If the router receives packets from the host, there is a problem with the host or the application. Consult the documentation for the host or application. If the router does receive packets from the host, use the **show running-config** privileged EXEC command to check the configuration of the router interface that first receives the packet from the host. **Step 3** Look for an **ip helper-address** *address* interface configuration command entry for that interface. Make sure that the specified address is correct (it should be the IP address of a server application such as a BOOTP server). If there is no command entry then no helper address is configured. **Step 4** If there is no IP helper address configured, or if the wrong address is specified, add or change the helper address using the **ip helper-address** *address* interface configuration command. For example, to configure the IP address 192.168.192.6 as the helper address on router Ethernet interface 0, enter the following commands: `C4500(config)#interface e0` `C4500(config-if)#ip helper-address 192.168.192.6`
UDP broadcasts being forwarded out nondefault ports	Specifying an IP helper address ensures only that broadcasts from a certain default set of UDP ports are forwarded. UDP broadcasts forwarded out other ports require further configuration. Enter an **ip forward-protocol udp** *port* global configuration command on the router for each applicable port. For example, to forward UDP broadcasts from port 200, enter the following command: `C4500(config)#ip forward-protocol udp 200` To allow forwarding of all UDP broadcasts, enter the following command: `C4500(config)#ip forward-protocol udp`

Possible Problem	Solution
UDP broadcast forwarding is disabled on specific UDP ports	**Step 1** Use the **show running-config** privileged EXEC command on the router and look for any **no ip forward-protocol udp** global configuration command entries. Such entries disable the forwarding of UDP traffic out specific ports. For example, entering the **no ip forward-protocol udp 53** global configuration command will disable the forwarding of all UDP traffic out port 53, which is the default port for DNS broadcasts. The following entry is shown in the configuration: `no ip forward-protocol udp domain` **Step 2** If UDP broadcasts are disabled at specific UDP ports, enter the **ip forward-protocol udp** *port* global configuration command (you can also specify a keyword, such as **domain**, rather than the port number). For example, to reenable DNS broadcasts, enter the following command: `C4500(config)#ip forward-protocol udp domain` To allow forwarding of BOOTP broadcasts, enter the following command: `C4500(config)#ip forward-protocol udp bootp` To allow forwarding of all UDP broadcasts, enter the following command: `C4500(config)#ip forward-protocol udp`
Access list or other filters are misconfigured	**Step 1** Use the **show running-config** command to check the configuration of each router in the path. See if there are access lists configured on the router. **Step 2** If there are access lists enabled on the router, disable them using the appropriate commands. For example, to disable input access list 10, enter the following command: `C4000(config-if)#no ip access-group 10 in` **Step 3** After disabling all access lists, determine if the BOOTP or other UDP broadcasts are forwarded normally. **Step 4** If broadcasts are forwarded normally, an access list is probably blocking traffic. **Step 5** To isolate the problem access list, enable access lists one at a time until broadcasts are no longer forwarded. **Step 6** Check the problem access list to see if it is filtering traffic from any UDP ports. If an access list denies specific UDP ports, make sure that it does not deny ports used to forward the broadcast traffic in question (such as UDP port 67 for BOOTP or port 68 for BOOTP replies). Enter explicit **permit** statements for those ports used to forward broadcasts that you want to have forwarded. **Step 7** If you altered an access list, enable the list to see if broadcasts are still forwarded normally. **Step 8** If problems persist, perform the preceding steps on routers in the path until broadcast traffic is forwarded correctly. For more information about misconfigured access lists, see the section “Misconfigured Access List Example” earlier in this chapter. For more information on configuring access lists, see the Cisco IOS configuration guides and command references.

TCP/IP: Poor Performance

Symptom: Performance for one or more network hosts is slow. Connections to servers take an excessive amount of time to establish.

Table 5-6 outlines the problems that might cause this symptom and describes solutions to those problems.

Table 5-6 TCP/IP: Poor Performance

Possible Problem	Solution
Misconfigured resolv.conf file on DNS client	Check the /etc/resolv.conf file on DNS clients. If the file is misconfigured, the client might wait until a query to one server times out before trying a second server, an NIS[1], or its host tables. This can cause excessive delays.
DNS is not set up for reverse lookups	If the DNS server is not configured to perform reverse lookups, reverse lookup attempts by end systems will time out. This can cause excessive delays for hosts attempting to establish connections. Consult your DNS software documentation for information on how to properly configure the DNS for reverse lookups.
DNS host table is incomplete	If the DNS host table is incomplete, reverse lookups will be unsuccessful, causing timeouts and therefore delays. **Step 1** At the UNIX prompt, enter the following command: `unix-host% host ip-address` where *ip-address* is the IP address of a server, router, or other network node. **Step 2** If the result of this command is "Host not found," but you can open the connection using the host's IP address rather than its name, then the host table might be incomplete. Add address-to-hostname mappings to the DNS host table for every host on the network.

1. NIS=Network Information Service

RIP/IGRP: Routes Missing from Routing Table

Symptom: Routes are missing from the routing table. Hosts on one network cannot access hosts on a different network. Error messages stating “host or destination unreachable” are generated.

The problem might be occurring in an internetwork running only RIP or IGRP, or a combination of the two.

Table 5-7 outlines the problems that might cause this symptom and describes solutions to those problems.

Table 5-7 RIP/IGRP: Routes Missing from Routing Table

Possible Problem	Solution
Misconfigured or missing **network** router configuration command	**Step 1** Use the **show running-config** privileged EXEC command to view the router configuration. **Step 2** Make sure that a **network** router configuration command is specified for every network to which a router interface belongs. For example, if the IP address of one interface is 192.168.52.42, and the IP address of another interface is 108.168.54.10, enter the following commands to enable RIP on the interfaces: `c4500(config)#router rip` `c4500(config-router)#network 192.168.52.0` `c4500(config-router)#network 108.168.0.0` Make sure the proper process IDs, addresses, and other variables are properly specified for the routing protocol you are using. For more information, consult the Cisco IOS configuration guides and command references.

Possible Problem	Solution
Misconfigured route filtering	**Step 1** Use the **show running-config** command to check suspect routers. **Step 2** See if any **distribute-list in** or **distribute-list out** router configuration commands are configured on the router. The **distribute-list in** command filters specific information in routing updates received by a router. The **distribute-list out** command prevents a router from including specific information in routing updates that it transmits. The information that is filtered is specified with an access list. **Step 3** If **distribute-list** commands are configured on the router, disable them using the **no** version of the command. For example, to disable an incoming filter that references access list 10, enter the following command: `C7500(config)#no distribute-list 10 in` **Step 4** After disabling all distribution lists on the router, use the **clear ip route** privileged EXEC command to clear the routing table. **Step 5** Determine if the routes appear in the routing table by using the **show ip route** privileged EXEC command. **Step 6** If routes appear properly in the routing table, the access list referenced by the **distribute-list** command is probably configured to deny certain updates. **Step 7** To isolate the problem list, enable distribution lists until routes stop appearing in the routing table. (You might have to use the **clear ip route** command after enabling each list.) **Step 8** Use the **show running-config** command and make sure that the problem list does not deny updates inappropriately. If the access list denies updates from specific addresses, make sure that it does not deny the address of a router from which routing updates should be received. Change the access list to allow the router to receive updates from the proper addresses. Remember that an implicit **deny any** ends every access-list. Configure explicit **permit** statements for those addresses from which the router should receive updates. **Step 9** If you altered an access list, enable the distribution list using the **distribute-list** command. Use the **clear ip route** command and check to see if the missing routing information appears in the routing table. **Step 10** If the routes appear, perform the preceding steps on all routers in the path until the routing information appears properly with all distribution lists enabled. For more information on configuring access lists, see the Cisco IOS configuration guides and command references.

Possible Problem	Solution
Subnet mask mismatch	Problems occur when two or more interfaces on the same major network have different subnet masks configured. **Step 1** Use the **show running-config** privileged EXEC command to view the configuration of each router in the major network. **Step 2** Use the **show ip interface** privileged EXEC command. Check the subnet mask specified for each interface. There is a subnet mask mismatch if two or more interfaces on the same major network have different subnet masks. **Step 3** If two interfaces on the same network have different subnet masks, you must change the subnet mask specification for one of the interfaces using the **ip address** *ip-address mask* interface configuration command (or use a classless routing protocol such as OSPF or Enhanced IGRP). For example, to configure Ethernet interface 1 with the IP address 192.168.52.46 using a subnet mask of 255.255.255.0, enter the following commands: `C4000(config)#interface e1` `c4000(config-if)#ip address 192.168.52.46 255.255.255.0` For more information about subnet masks, see the section "Host and Router Subnet Mask Mismatch Example" later in this chapter.
Missing **default-metric** command	This problem is restricted to environments in which route redistribution is being performed between autonomous systems or between multiple routing protocols. **Step 1** Use the **show running-config** privileged EXEC command on suspect routers. Look for **default-metric** router configuration command entries. This command assigns default metric values to redistributed routes. **Step 2** IGRP requires a default-metric parameter to redistribute routes. If you are running IGRP, define the default metrics for redistributed routes using the **default-metric** router configuration command. The following example shows a configuration that redistributes RIP routes and assigns them IGRP metrics with values as follows: bandwidth = 1000, delay = 100, reliability = 250, loading = 100, and mtu = 1500. `router igrp 109` `network 131.108.0.0` `redistribute rip` `default-metric 1000 100 250 100 1500` **Step 3** If you are running RIP, you do not have to configure a default metric in order to redistribute routes. By default, the metric assigned to all routes redistributed into RIP is 1. However, this value can be changed using the **default-metric** command. If a **default-metric** statement that is applied to RIP appears in the configuration, make sure that the metric value it assigns will not adversely affect network performance. If you are unsure, restore the default value for the routing metric using the **no default-metric** router configuration command. For more information on the **default-metric** router configuration command, see the Cisco IOS configuration guides and command references.

Possible Problem	Solution
Routes are not being redistributed properly between autonomous systems or between routing protocols	This problem is restricted to environments in which route redistribution is being performed between autonomous systems, or between multiple routing protocols. **Step 1** Use the **show running-config** privileged EXEC command on routers that border multiple networks running different routing protocols. **Step 2** Examine the **router** global configuration command entries for the enabled routing protocols. **Step 3** If the router is running IGRP only, check whether the autonomous system designated for all connected networks is the same. Routes are not automatically redistributed between different autonomous systems. If the **router igrp** commands indicate different autonomous systems, route redistribution must be manually configured using the **redistribute** router configuration command. For example, to redistribute routes between IGRP autonomous system 71 (network 15.0.0.0) and IGRP autonomous system 109 (network 192.31.7.0), enter the following commands: `C7010(config)#router igrp 71` `C7010(config-router)#redistribute igrp 109` `C7010(config-router)#distribute-list 3 out igrp 109` `C7010(config-router)#access-list 3 permit 192.31.7.0` `C7010(config)#router igrp 109` `C7010(config-router)#redistribute igrp 71` `C7010(config-router)#distribute-list 5 out igrp 71` `C7010(config-router)#access-list 5 permit 15.0.0.0` **Step 4** If the router is running multiple routing protocols, look for a **redistribute** router configuration command entry. Make sure that routing information is being properly exchanged between protocols. For example, to redistribute routes between RIP (running in network 15.0.0.0) and IGRP autonomous system 109 (network 128.1.0.0), enter the following commands: `C7010(config)#router igrp 109` `C7010(config-router)#network 128.1.0.0` `C7010(config-router)#redistribute rip` `C7010(config-router)#default-metric 10000 100 255 1 1500` `C7010(config-router)#distribute-list 10 out rip` `C7010(config-router)#access-list 10 permit 15.0.0.0` **Step 5** If you want static routes to be redistributed between autonomous systems or between two different routing protocols, use the **redistribute static** router configuration command. For example, to redistribute static routes IGRP autonomous systems, add the following command to the configuration: `C7010(config-router)#redistribute static` For more information on using the **redistribute** router configuration command, see the Cisco IOS configuration guides and command references.

Host and Router Subnet Mask Mismatch Example

In classful IP networks, every router and host in the same major network should share a common subnet mask. If there are disagreements on the length of the subnet mask, packets will not be routed correctly.

Table 5-8 shows how a UNIX host and a router will interpret an IP address differently if they have different subnet masks specified for the same major network.

Table 5-8 Host and Router Subnet Mask Mismatch Example

Routing Information	Host Value	Router Value
Destination IP address	192.31.7.49	192.31.7.49
Subnet mask	255.255.255.240	255.255.255.224
Interpreted address	Subnet address 48, host 1	Subnet address 32, host 17

The host interprets the IP address 192.31.7.49 as being Host 1 on the third subnet (subnet address 48). However, because it is using a different subnet mask, the router interprets the address as being to Host 17 on the first subnet (subnet address 32). Depending on the network topology and the router configuration, packets destined for IP address 192.31.7.49 might be sent to the wrong destination host, sent out the wrong interface, or dropped altogether.

OSPF: Routers Not Establishing Neighbors

Symptom: OSPF routers are not establishing neighbor relationships properly. The result is that routing information is not exchanged between routers.

Table 5-9 outlines the problems that might cause this symptom and describes solutions to those problems.

Table 5-9 OSPF: Routers Not Establishing Neighbors

Possible Problem	Solution
Misconfigured or missing **network** router configuration command	**Step 1** Use the **show ip ospf interfaces** EXEC command to see which interfaces have OSPF enabled. **Step 2** If the output indicates that an interface that should be running OSPF is not doing so, use the **show running-config** privileged EXEC command to view the router configuration. **Step 3** Make sure that **network** router configuration commands are specified for each interface on which OSPF should run. For example, if the IP address of Ethernet interface 0 is 192.168.52.42 with a subnet mask of 255.255.255.0, enter the following commands to enable OSPF on the interface: `c4500(config)#router ospf 100` `c4500(config-router)#network 192.168.52.0 0.0.0.255 area 0` Make sure the proper process IDs, addresses, wildcard masks, and other variables are properly specified. **Note:** There is no correlation between OSPF wildcard masks (used in OSPF **network** commands) and the subnet mask configured as part of an interface IP address. **Step 4** Check other OSPF routers on the network using the preceding steps. Make sure that OSPF is configured properly on all neighboring routers so that neighbor relationships can be established.

Possible Problem	Solution
Mismatched Hello or dead timers, E-bits (set for stub areas), area IDs, authentication types, or network masks	The values set for the Hello timer and dead timer intervals, E-bits (this bit is set if the router is configured in a stub area), area IDs, authentication types, and network masks should all be the same throughout an OSPF area and in some cases the entire OSPF network. **Step 1** Use the **show ip ospf neighbor** privileged EXEC command to identify the OSPF neighbors of each router. **Step 2** If the output does not list an expected neighbor, use the **show ip ospf interface** privileged EXEC command on the router and its expected neighbor. Examine the Hello and dead timer interval values configured on OSPF interfaces. `C7010#show ip ospf interface` `[...]` `Timer intervals configured, Hello 12, Dead 48, Wait 40, Retransmit 5` **Step 3** Compare the values configured for the timers on each router. If there is a mismatch, reconfigure the timer values so that they are the same on the router and its neighbor. For example, to change the Hello timer interval to 10 on Ethernet interface 0/1, enter the following commands: `C7010(config)#interface e0/1` `C7010(config-if)#ip ospf hello-interval 10` **Step 4** Use the **debug ip ospf adj** privileged EXEC command. Check the output for mismatched values. In the following example, there is a network mask mismatch. The mask received from router 141.108.10.3 is 255.255.255.0, and the mask configured on the router C4500 is 255.255.255.252: `C4500#debug ip ospf adj` `OSPF: Mismatched hello parameters from 141.108.10.3` `Dead R 40 C 40, Hello R 10 C 10 Mask R 255.255.255.0 C 255.255.255.252` **Step 5** If mismatches are indicated in the debug output, try to resolve the mismatch. For detailed information about configuring OSPF, see the Cisco IOS *Network Protocols Configuration Guide, Part 1*. **Step 6** Perform the same types of steps for all of these parameters. Check that all routers in an area have the same area ID, whether all routers in the area are configured as stub routers, whether the same authentication type is configured for all routers, and so forth. For information on configuring these parameters, consult the Cisco IOS *Network Protocols Configuration Guide, Part 1*. **Note:** Timer values are extremely important when Cisco routers interoperate with routers from other vendors.

Possible Problem	Solution
Access list is misconfigured	**Step 1** Use the **show access-list** privileged EXEC command on suspect routers to see if there are IP access lists configured on the router. **Step 2** If there are IP access lists enabled on the router, disable them using the appropriate commands. For example, to disable input access list 10, use the following command: `C4000(config-if)#no ip access-group 10 in` **Step 3** After disabling all access lists on the router, determine if the router is able to establish neighbor relationships normally. Use the **show ip ospf neighbor** privileged EXEC command. If the proper neighbor relationships have been established, an access list is probably filtering OSPF hello packets. **Step 4** To isolate the problem access list, enable access lists one at a time until the router cannot establish neighbors (use the **clear ip ospf neighbors** privileged EXEC command to force the router to clear the neighbor table). **Step 5** Check the access list to see if it is filtering traffic from port 89, the port used by OSPF. Remember that every access list ends with an implicit **deny any** statement. If an access list denies OSPF traffic, enter an explicit **permit** statement for port 89 to ensure that neighbor relationships can be established properly. (You can also use the **ospf** keyword when configuring the access list.) For example, to configure input access list 101 to allow OSPF traffic to pass, enter the following commands on the router: `C4500(config)#access-list 101 permit ospf any any` **Step 6** If you altered an access list, enable the list and enter the **clear ip ospf neighbors** privileged EXEC command. Then enter the **show ip ospf neighbor** command to see if neighbor relationships are established normally. **Step 7** If the router is establishing neighbors, perform the preceding steps for other routers in the path until all access lists are enabled and the router can still establish neighbors normally. For more information on configuring access lists, see the Cisco IOS configuration guides.
Virtual link and stub area configuration mismatch	**Step 1** A virtual link cannot be configured across a stub area. Check router configurations for routers configured both as part of a stub area and as an ABR[1] that is part of a virtual link. Use the **show running-config** privileged EXEC command and look for command entries that are similar to the following: `area 2 stub` `area 2 virtual-link 192.169.100.10` **Step 2** If both of these commands are present, there is a misconfiguration. Remove one of the commands (using the **no** form of the command) to resolve the misconfiguration.

1. ABR=area border router

OSPF: Routes Missing from Routing Table

Symptom: OSPF routes and networks are not being advertised to other routers. Routers in one area are not receiving routing information for other areas. Some hosts cannot communicate with hosts in other areas, and routing table information is incomplete.

Table 5-10 outlines the problems that might cause this symptom and describes solutions to those problems.

Table 5-10 OSPF: Routes Missing from Routing Table

Possible Problem	Solution
OSPF routers not establishing neighbors	Follow the procedures outlined in the section "OSPF: Routers Not Establishing Neighbors" earlier in this chapter.
Routing information from IGRP or RIP is not redistributed correctly into OSPF	**Step 1** Check the router configuration using the **show running-config** privileged EXEC command. **Step 2** Look for a **redistribute** router configuration command entry. Make sure that redistribution is configured and that the **subnets** keyword is used with the command. The **subnets** keyword must be included when IGRP or RIP is redistributed into OSPF; otherwise, only major routes (not subnet routes) are redistributed. **Step 3** If the **redistribute** command is not present, or if the **subnets** keyword is not specified, add or change the configuration using the following commands: `C7000(config)#router ospf 100` `C7000(config)#redistribute ospf subnets`
No ABR is configured in an area, isolating that area from the OSPF backbone	**Step 1** Use the **show running-config** privileged EXEC command on OSPF routers to verify that at least one ABR exists for the area. ABRs must belong to area 0, the OSPF backbone, as well as to another area. Look for **network** statements that indicate that the router is part of area 0. **Step 2** If no ABR exists in an area, configure one where appropriate. Use the **network** router configuration command. For example, to configure OSPF process 100 to participate in the OSPF backbone area, enter the following commands: `C4500(config)#router ospf 100` `C4500(config-router)#network 192.21.3.7 0.0.0.255 area 0`

Possible Problem	Solution
Interface network type mismatch on Frame Relay WAN	In an OSPF Frame Relay environment, if one end of the link is a multipoint interface and the other end is a point-to-point interface, by default the multipoint interface will advertise the link as a non-broadcast network and the point-to-point interface will advertise the link as a point-to-point network. This creates a conflict in the link-state database and can prevent routing information from being learned properly. **Step 1** Check each router interface on each side of the link to see if the network types are mismatched. Use the **show ip ospf interface** privileged EXEC command to check the network type configured for the interface. Following is an example of the output from the **show ip ospf interface** command: `Ethernet0 is up, line protocol is up` `Internet Address 192.168.52.14 255.255.255.0, Area 0` `Process ID 1, Router ID 192.168.52.14, Network Type BROADCAST, Cost: 10` `[...]` In this example, the network type is broadcast. **Step 2** Change the point-to-point interface to a multipoint interface by configuring subinterfaces, or change the network type of the point-to-point interface to broadcast using the **ip ospf network broadcast** interface configuration command. For information on configuring subinterfaces, see the Cisco IOS configuration guides.
Area is configured as a stub area	Route redistribution is not possible in OSPF stub areas. No external routes are advertised into a stub area, and if the **area** *area-id* **stub no-summary** router configuration command is used, no summary routes (inter-area routes) will be advertised into the stub area. **Step 1** If you want summary routes to be advertised into the stub area, but you do not see them in the routing table, use the **show running-config** privileged EXEC command to view the router configuration. **Step 2** Look for an **area** *area-id* **stub no-summary** command entry. If this command is present, disable it by entering the following commands: `C4500(config)#router ospf 100` `C4500(config-router)#no area 1 stub no-summary` This disables the **no-summary** keyword and keeps the router configured as a stub. **Step 3** To advertise external routes into the area, you must configure the area as a non-stub. Make certain that all routers in the area are reconfigured as non-stub routers.

Possible Problem	Solution
Misconfigured route filtering	**Step 1** Use the **show running-config** command to check suspect routers. **Step 2** See if there are any **distribute-list in** or **distribute-list out** router configuration commands configured on the router. The **distribute-list in** command prevents specific information learned in LSAs[1] from being included in the OSPF routing table. The **distribute-list out** command prevents a router from including specific information in routing updates that it transmits. However, in OSPF, **distribute-list out** can be configured *only* on an ASBR[2] to filter external routes. **Note:** Although **distribute-list** commands prevent specific information from being included in the OSPF routing table, information about those networks is contained in the link-state database and is flooded through the network in LSAs. This means that downstream routers will include that information in their routing tables unless they too filter those routes from the routing table. **Step 3** If **distribute-list** commands are configured on the router, disable them using the **no** version of the command. For example, to disable an incoming filter that references access list 10, enter the following command: `C7500(config)#no distribute-list 10 in` **Step 4** After disabling all distribution lists, use the **clear ip route** privileged EXEC command to clear the routing table. **Step 5** Determine if the routes appear in the routing table by using the **show ip route** privileged EXEC command. If routes appear properly in the routing table, the access list referenced by the **distribute-list** command is probably configured to deny certain updates. **Step 6** To isolate the problem list, enable distribution lists one at a time until the routes no longer appear in the table. **Step 7** Use the **show running-config** command and check the access list to make sure it does not deny updates inappropriately. If the access list denies updates from specific addresses, make sure that it does not deny the address of a router from which routing updates should be received. Change the access list to allow the router to receive updates from the proper addresses. Remember that an implicit **deny any** ends every access-list. Configure explicit **permit** statements for those addresses from which the router should receive updates. **Step 8** If you altered an access list, enable the distribution list using the **distribute-list** command. Use the **clear ip route** command and check to see if the missing routing information appears in the routing table. **Step 9** If the routes appear in the routing table, perform the preceding steps on every router in the path until all distribution lists are enabled and routing information appears properly in the routing table. For more information on configuring access lists, see the Cisco IOS configuration guides.

Possible Problem	Solution
Virtual link is misconfigured	**Step 1** Check the configuration of the routers at each end of the virtual link using the **show running-config** privileged EXEC command. Look for **area** *area-id* **virtual-link** *router-id* router configuration command entries. These commands are used to configure the virtual link. **Step 1** Use the **show ip ospf** EXEC command to find the router ID (IP address) of the routers. **Step 2** Add the **area** *area-id* **virtual-link** *router-id* command if it is missing, or modify it if it is incorrect. Make sure that the proper area ID and router ID (IP address) are specified. The routers at each end of the virtual link must point to one another across the transit area. For example, in the network shown in Figure 5-3, a virtual link from Router B to Router A is created across the transit area, Area 1. The following commands are entered on Router A: `C4500(config)#router ospf 250` `C4500(config-router)#network 121.10.0.0 0.0.255.255 area 0` `C4500(config-router)#network 169.192.56.0 0.0.0.255 area 0` `C4500(config-router)#area 1 virtual-link 121.10.100.46` On Router B, the following commands are used: `C4000(config)#router ospf 250` `C4000(config-router)#network 121.10.0.0 0.0.255.255 area 0` `C4000(config-router)#network 108.31.0.0 0.0.255.255 area 2` `C4000(config-router)#area 1 virtual-link 121.10.1.1`

1. LSA=link state advertisement
2. ASBR=autonomous system border router

Figure 5-3 OSPF Virtual Link Example

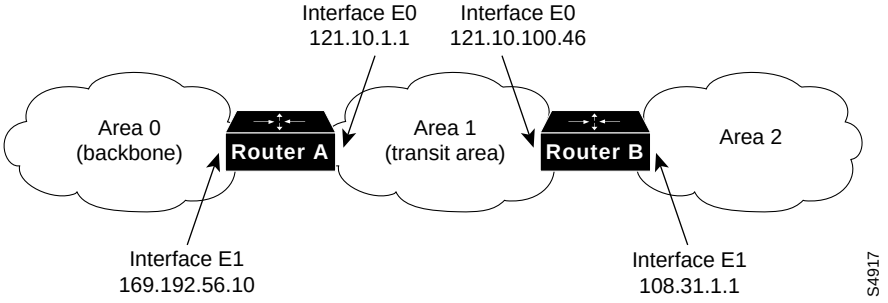

IP Enhanced IGRP: Routers Not Establishing Neighbors

Symptom: Enhanced IGRP routers are not establishing neighbor relationships properly. Routing information is not distributed to routers.

Table 5-11 outlines the problems that might cause this symptom and describes solutions to those problems.

Table 5-11 IP Enhanced IGRP: Routers Not Establishing Neighbors

Possible Problem	Solution
Misconfigured or missing **network** router configuration command	**Step 1** Use the **show ip eigrp neighbors** EXEC command on an Enhanced IGRP router. Make sure that all directly connected Enhanced IGRP routers appear in the output. **Step 2** If some connected routers are not shown in the output, use the **show running-config** privileged EXEC command to view the configuration of the suspect routers. **Step 3** Make sure that a **network** router configuration command is specified for every network to which a router interface belongs. For example, if the IP address of Ethernet interface 0 is 192.168.52.42, enter the following commands to enable Enhanced IGRP on the interface: `c4500(config)#router eigrp 100` `c4500(config-router)#network 192.168.52.0`
Mismatched autonomous system number specification	**Step 1** View the router configuration using the **show running-config** privileged EXEC command on each router in the autonomous system. **Step 2** Check the **router eigrp** global configuration commands to make sure that all routers which should be communicating are in the same autonomous system. Only Enhanced IGRP routers in the same autonomous system will form neighbor relationships and thus exchange routing information.
Access list is misconfigured	**Step 1** Enable the **debug ip packet** and **debug eigrp packets** privileged EXEC commands. The former command indicates whether IP packets are being sent and received, and whether there are encapsulation problems. The latter command indicates whether Enhanced IGRP hello packets are being sent and received properly. **Caution:** These **debug** commands can use considerable CPU cycles on the router. Do not enable them if your network is already heavily congested. **Step 2** If a router appears to be sending IP and Enhanced IGRP packets correctly, but a connected router does not receive them, check the configuration of the connected router for access lists that might be filtering out packets. **Step 3** Disable all access lists enabled on the router using the **no ip access-group** *access-list-number* **in** interface configuration command. **Step 4** Monitor the output from the **debug ip packet** and **debug eigrp packets** commands. Determine if packets are now being received normally. **Step 5** If packets are received normally, an access list is probably filtering packets. To isolate the problem list, enable access lists one at a time until packets are no longer forwarded. **Step 6** Check the access list to see if it is filtering traffic from the source router. If it is, alter the access list to allow the traffic to pass. Enter explicit **permit** statements for traffic that you want the router to forward normally. **Step 7** Enable the altered access list with the **ip access-group** command to see if packets continue to pass normally. **Step 8** If packets pass normally, perform the preceding steps on any other routers in the path until all access lists are enabled and packets are forwarded properly.

IP Enhanced IGRP: Routes Missing from Routing Table

Symptom: Routes are missing from the routing table of routers running Enhanced IGRP. Hosts on one network cannot access hosts on a different network. Hosts on the same network might or might not be able to communicate. The problem might occur in internetworks running only Enhanced IGRP, or in an internetwork running Enhanced IGRP and another routing protocol.

Table 5-12 outlines the problems that might cause this symptom and describes solutions to those problems.

Table 5-12 IP Enhanced IGRP: Routes Missing from Routing Table

Possible Problem	Solution
Routers not establishing neighbors	For information on troubleshooting this problem, see the section "IP Enhanced IGRP: Routers Not Establishing Neighbors," earlier in this chapter.
Routes are not redistributed between different autonomous systems	Routes are not automatically redistributed between different Enhanced IGRP autonomous systems. **Step 1** Use the **show running-config** privileged EXEC command on routers bordering multiple autonomous systems. **Step 2** If multiple autonomous systems are configured on the router (indicated by multiple **router eigrp** global configuration command entries), make sure that route redistribution is manually configured using the **redistribute** router configuration command. For example, if the router belongs to autonomous system 100 and autonomous system 200, enter the following commands to redistribute Enhanced IGRP routes between the two autonomous systems: `C2509(config)#router eigrp 100` `C2509(config-router)#redistribute eigrp 200` `C2509(config-router)#exit` `C2509(config)#router eigrp 200` `C2509(config-router)#redistribute eigrp 100` **Step 3** If you want static routes to be redistributed, you must use the **redistribute static** router configuration command. For more information on using the **redistribute** router configuration command, see the Cisco IOS configuration guides and command references.
Routes are not being redistributed between different routing protocols	**Step 1** Use the **show running-config** privileged EXEC command on routers that border networks running different routing protocols. **Step 2** Look for a **redistribute** router configuration command entry. Make sure that routing information is being properly exchanged between protocols. For example, to redistribute routes between IGRP autonomous system 500 and Enhanced IGRP autonomous system 200, enter the following commands: `C2509(config)#router igrp 500` `C2509(config-router)#redistribute eigrp 200` `C2509(config-router)#exit` `C2509(config)#router eigrp 200` `C2509(config-router)#redistribute igrp 500` **Step 3** To redistribute static routes, you must use the **redistribute static** router configuration command. For more information on using the **redistribute** router configuration command, see the Cisco IOS configuration guides and command references.

Possible Problem	Solution
Hello interval or hold-time value mismatch	**Step 1** Use the **show running-config** privileged EXEC command on all routers in the network. **Step 2** Look for **ip hello-interval eigrp** and **ip hold-time eigrp** interface configuration command entries. The values configured by these commands should be the same for all IP routers on the network. At minimum, backbone routers should be configured with the same hello interval and hold-time values. **Step 3** If there are routers with mismatched hello interval or hold-time values, reconfigure them to bring them into conformance with the rest of the routers on the network. You can return these timer values to their defaults by using the **no ip hello-interval eigrp** and the **no ip hold-time interval eigrp** interface configuration commands.
Default routing metrics are incorrectly configured	**Step 1** Use the **show running-config** privileged EXEC command on suspect routers. Look for **default-metric** router configuration command entries. This command changes the default metric values assigned to redistributed routes. **Step 2** If a **default-metric** statement appears in the configuration, examine the values that it defines. Be certain that these values will reliably and accurately translate routing metrics between the routing protocols implemented on your network. To restore the default values for the routing metrics, use the **no default-metric** router configuration command for the appropriate routing protocol. For more information on the IP Enhanced IGRP **default-metric** router configuration command, see the Cisco IOS configuration guides.

IP Enhanced IGRP: Router Stuck in Active Mode

Symptom: An IP Enhanced IGRP router is stuck in Active mode. Multiple "Stuck-in-Active" messages are sent to the console:

```
%DUAL-3-SIA: Route 198.169.52.51 Stuck-in-Active
```

For a more detailed explanation of Enhanced IGRP Active mode, see the section "Enhanced IGRP and Active/Passive Modes" later in this chapter.

Note Occasional messages of this type are *not* a cause for concern. This is how an Enhanced IGRP router recovers if it does not receive replies to its queries from all of its neighbors. However, if these error messages occur frequently, you should investigate the problem.

Table 5-13 outlines the problems that might cause this symptom and describes solutions to those problems.

Table 5-13 IP Enhanced IGRP: Router Stuck in Active Mode

Possible Problem	Solution
Active timer value is misconfigured	**Step 1** Check the configuration of each Enhanced IGRP router using the **show running-config** privileged EXEC command. **Step 2** Look for the **timers active-time** router configuration command entry associated with the **router eigrp** global configuration command entry. The active timer determines the maximum period of time that an Enhanced IGRP router will wait for replies to its queries. If the active timer value is set too low, there might not be enough time for all of the neighboring routers to send their replies to the active router. **Step 3** Make sure that the value set by the **timers active-time** command is consistent among routers in the same autonomous system. A value of 3 (3 minutes, which is the default value) is recommended in order to allow all Enhanced IGRP neighbors to reply to queries.
Interface or other hardware problem	**Step 1** Use the **show ip eigrp neighbors** EXEC command and examine the Uptime and Q Cnt (queue count) fields in the output. If the uptime counter is continually resetting or if the queue count is consistently high, there might be a hardware problem. **Step 2** Determine where the problem is occurring by looking at the output of the Stuck in Active error message, which will indicate the direction in which the problem node is located. **Step 3** Make sure the suspect router still works. Check the interfaces on the suspect router. For more information, see the "Troubleshooting Hardware and Booting Problems" chapter.
Flapping route	**Step 1** If there is a flapping serial route (caused by heavy traffic load), queries and replies might not be forwarded reliably. Route flapping caused by heavy traffic on a serial link can cause queries and replies to be lost, resulting in the active timer timing out. **Step 2** Increase the bandwidth of the link. For more information, see the "Troubleshooting Serial Line Problems" chapter.

Enhanced IGRP and Active/Passive Modes

An Enhanced IGRP router can be in either Passive or Active mode. A router is said to be passive for a network when it has an established path to that network in its routing table.

If the Enhanced IGRP router loses the connection to a network (for example, Network A), it becomes active for that network. The router sends out queries to all of its neighbors in order to find a new route to Network A. The router remains in active mode until it has either received replies from *all* of its neighbors or until the active timer, which determines the maximum period of time a router will stay active, expires.

If the router receives a reply from each of its neighbors, it computes the new next hop to Network A and becomes passive for that network. However, if the active timer expires before all of its neighbors reply, the router removes from its neighbor table any neighbors that did not reply, again enters active mode, and sends a “Stuck-in-Active” message to the console.

BGP: Routes Missing from Routing Table

Symptom: BGP routers and networks are not advertised to other routers. Routers do not receive routing information from other routers. Some hosts cannot communicate with hosts in other areas, and routing table information is incomplete.

Table 5-14 outlines the problems that might cause this symptom and describes solutions to those problems.

Table 5-14 BGP: Routes Missing from Routing Table

Possible Problem	Solution
BGP routers not advertising routes	If BGP routers are not advertising routes properly, routing information might not appear in the routing table. For information on troubleshooting this problem, see the section "BGP: Routers Not Advertising Routes," later in this chapter.
Missing **neighbor remote-as** command	The **neighbor remote-as** router configuration command is used to add entries to the BGP neighbor table. **Step 1** Check local and remote routers and make sure the specified autonomous system numbers and neighbors are correct. **Step 2** Make sure any route filters that are enabled are not misconfigured.
Access list is misconfigured	**Step 1** Use the **show access-list** privileged EXEC command on suspect routers to see if there are access lists configured on the router. **Step 2** If there are access lists enabled on the router, disable them using the appropriate commands. For example, to disable input access list 10, use the following command: `C4000(config)#no ip access-group 10 in` **Step 3** After disabling all access lists on the router, determine if the missing routing information is now appearing in routing tables. **Step 4** If the information is now appearing, it is likely that an access list is filtering traffic. To isolate the problem access list, enable access lists one at a time until the routing information is no longer appearing in the routing table. **Step 5** Check the access list to see if it is filtering traffic from specific TCP ports. If an access list denies specific TCP ports, make sure that it does not deny TCP port 179, the port used by BGP. Enter an explicit **permit** statement for port 179 to ensure that BGP traffic is forwarded normally. **Step 6** If you altered an access list, enable the list to see if routing information can still pass normally. **Step 7** If routing information is no longer missing, perform the preceding steps on any other routers in the path until all access lists are enabled and routing information appears in the appropriate routing tables. For more information on configuring access lists, see the Cisco IOS configuration guides.

BGP: Routers Not Advertising Routes

Symptom: BGP routers are not advertising routes. Routing updates from a BGP router do not contain information about certain network destinations that should be advertised.

Table 5-15 outlines the problems that might cause this symptom and describes solutions to those problems.

Table 5-15 BGP: Routers Not Advertising Routes

Possible Problem	Solution
Missing **network** router configuration command	**Step 1** Use the **show running-config** privileged EXEC command to view the router configuration. **Step 2** Make sure that a **network** router configuration commands is specified for every network that the BGP router should advertise (these networks need not be directly connected). For example, if you want the BGP router to advertise networks 192.168.52.0 and 108.168.0.0, enter the following commands to have the router include those networks in its routing updates: `c4500(config)#router bgp 100` `c4500(config-router)#network 192.168.52.0` `c4500(config-router)#network 108.168.0.0`
Interior gateway protocol (such as RIP, IGRP, OSPF, and so on) routing problem	**Step 1** Check for other routing protocol problems to be sure that BGP is getting routing information from any interior gateway protocols running in the internetwork. For example, if there is a problem with RIP routing it might affect the operation of BGP. BGP routers might not have any information about certain networks, making it impossible to advertise routing information about certain networks configured in BGP. **Step 2** Isolate and troubleshoot interior gateway protocol problems before troubleshooting BGP. See the appropriate sections in this chapter for information specific to the protocols you are running. As a workaround, you can configure static BGP routes, but routing will not be dynamic in this case.
Misconfigured **aggregate-address** command	The **aggregate-address** router configuration command allows BGP to specify a summary address for one or more specific network addresses. For example, to summarize the addresses 195.10.20.0 and 195.10.130.0, use the aggregate address 195.10.0.0. Problems can occur under the following circumstances: • The aggregate address summarizes addresses that are not in the router's BGP routing table In this case, a router is advertising networks to which it does not have a BGP route. For example, a router is configured with the aggregate address 195.10.0.0 summarizing networks 195.10.20.0 and 195.10.130.0. However, network 195.10.192.0 is in another autonomous system that is inaccessible through the router. Traffic destined for network 195.10.192.0 will be forwarded to the router because it is incorrectly advertising a route to that network (via the aggregate address). • There are no individual networks configured (using the **network** router configuration command) or routes in the BGP routing table to which the aggregate address refers. **Step 1** Use the **show running-config** privileged EXEC command to view the router configuration. Look for an **aggregate-address** command entry associated with the **router bgp** global configuration command. **Step 2** Use the **show ip bgp** privileged EXEC command to view the addresses in the BGP routing table. **Step 3** Make sure that the addresses summarized by the **aggregate-address** command are all present in the BGP routing table.

HSRP: Hosts Cannot Reach Remote Networks

Symptom: Hosts cannot reach hosts on remote networks. Routers in the network are running HSRP.

Table 5-16 outlines the problems that might cause this symptom and describes solutions to those problems.

Table 5-16 HSRP: Hosts Cannot Reach Remote Networks

Possible Problem	Solution
Default gateway is not specified or is incorrectly specified on local or remote hosts	**Step 1** Determine whether local and remote hosts have a default gateway specification. Use the following UNIX command: `host% netstat -rn` Check the output of this command for a default gateway specification. **Step 2** In a network running HSRP, hosts must use the hot standby IP address as their default gateway specification. Use the **show standby** privileged EXEC command to check the current hot standby IP address. You can change or add a default gateway using the following UNIX command at the host: `host% route add default address 1` where *address* is the IP address of the default gateway (the router local to the host). The value 1 indicates that the specified gateway is one hop away. You might need to reboot the host for this change to take effect. **Step 3** It is recommended that you specify a default gateway as part of the boot process. Specify the IP address of the gateway in the following UNIX host file: */etc/defaultrouter* This filename might be different on your UNIX system. If you are working with a PC or a Macintosh, consult the accompanying documentation to determine how to set the default gateway.

Possible Problem	Solution
HSRP is not configured or is misconfigured	**Step 1** Try to **ping** the hot standby IP address. If the **ping** is unsuccessful, proceed to Step 2. If the **ping** is successful, proceed to Step 4. **Step 2** Use the **show standby** privileged EXEC command to see information about the HSRP configuration. If the command does not return any output, HSRP is not configured on the router interface. **Step 3** If HSRP is not configured, configure it on the routers that you want to belong to the hot standby group. For example, to configure a router as the active hot standby router with hot standby address 192.192.192.3, enter the following commands: `C4500(config)#interface e0` `C4500(config-if)#standby ip 192.192.192.3` `C4500(config-if)#standby priority 110` `C4500(config-if)#standby preempt` To configure a router as the backup hot standby router, enter the following commands: `C4500(config)#interface e0` `C4500(config-if)#standby ip 192.192.192.3` **Step 4** If the backup hot standby router is misconfigured and the active router fails, the backup router might not go active. One potential misconfiguration is a missing hot standby address in the backup router. A router can be configured successfully as a hot standby router simply by entering the following commands: `C4500(config)#interface e0` `C4500(config-if)#standby ip` That is, you do not have to include the hot standby IP address in the **standby ip** command. As long as one hot standby router has the hot standby IP address in its configuration, every other hot standby router will learn the address from that router. However, if only one router has the hot standby address configured, and that router fails, other hot standby routers will not know the hot standby address and HSRP will not work. Be sure that at least two hot standby routers have the hot standby address in their configuration.
No routes in active hot standby router	If HSRP appears to be configured correctly, but connectivity fails, make sure that your other routing protocols are working correctly. If your other routing protocols are not advertising routes correctly, hot standby routers will have incomplete or empty routing tables and traffic will not be forwarded correctly. Follow the troubleshooting procedures outlined in this chapter to ensure that your other routing protocols work correctly.