

UDP Broadcast Flooding

A *broadcast* is a data packet that is destined for multiple hosts. Broadcasts can occur at the data link layer and the network layer. Data-link broadcasts are sent to all hosts attached to a particular physical network. Network layer broadcasts are sent to all hosts attached to a particular logical network. The Transmission Control Protocol/Internet Protocol (TCP/IP) supports the following types of broadcast packets:

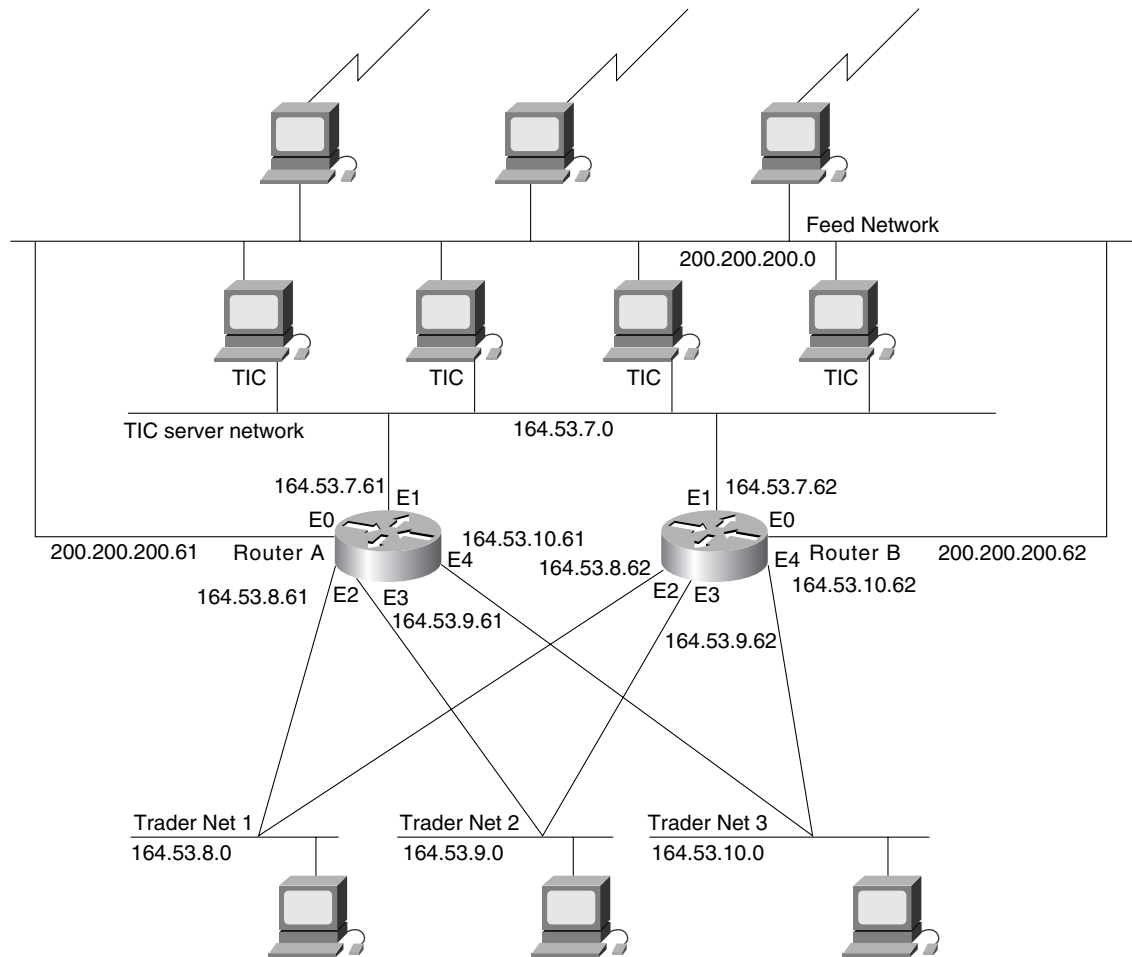
- *All ones*—By setting the broadcast address to all ones (255.255.255.255), all hosts on the network receive the broadcast.
- *Network*—By setting the broadcast address to a specific network number in the network portion of the IP address and setting all ones in the host portion of the broadcast address, all hosts on the specified network receive the broadcast. For example, when a broadcast packet is sent with the broadcast address of 131.108.255.255, all hosts on network number 131.108 receive the broadcast.
- *Subnet*—By setting the broadcast address to a specific network number and a specific subnet number, all hosts on the specified subnet receive the broadcast. For example, when a broadcast packet is set with the broadcast address of 131.108.4.255, all hosts on subnet 4 of network 131.108 receive the broadcast.

Because broadcasts are recognized by all hosts, a significant goal of router configuration is to control unnecessary proliferation of broadcast packets. Cisco routers support two kinds of broadcasts: *directed* and *flooded*. A directed broadcast is a packet sent to a specific network or series of networks, whereas a flooded broadcast is a packet sent to every network. In IP internetworks, most broadcasts take the form of User Datagram Protocol (UDP) broadcasts.

Although current IP implementations use a broadcast address of all ones, the first IP implementations used a broadcast address of all zeros. Many of the early implementations do not recognize broadcast addresses of all ones and fail to respond to the broadcast correctly. Other early implementations forward broadcasts of all ones, which causes a serious network overload known as a *broadcast storm*. Implementations that exhibit these problems include systems based on versions of BSD UNIX prior to Version 4.3.

In the brokerage community, applications use UDP broadcasts to transport market data to the desktops of traders on the trading floor. This case study gives examples of how brokerages have implemented both directed and flooding broadcast schemes in an environment that consists of Cisco routers and Sun workstations. Figure 6-1 illustrates a typical topology. Note that the addresses in this network use a 10-bit netmask of 255.255.255.192.

Figure 6-1 Topology that requires UDP broadcast forwarding.



In Figure 6-1, UDP broadcasts must be forwarded from a source segment (Feed network) to many destination segments that are connected redundantly. Financial market data, provided, for example, by Reuters, enters the network through the Sun workstations connected to the Feed network and is disseminated to the TIC servers. The TIC servers are Sun workstations running Teknekron Information Cluster software. The Sun workstations on the trader networks subscribe to the TIC servers for the delivery of certain market data, which the TIC servers deliver by means of UDP broadcasts. The two routers in this network provide redundancy so that if one router becomes unavailable, the other router can assume the load of the failed router without intervention from an operator. The connection between each router and the Feed network is for network administration purposes only and does not carry user traffic.

Two different approaches can be used to configure Cisco routers for forwarding UDP broadcast traffic: IP helper addressing and UDP flooding. This case study analyzes the advantages and disadvantages of each approach.

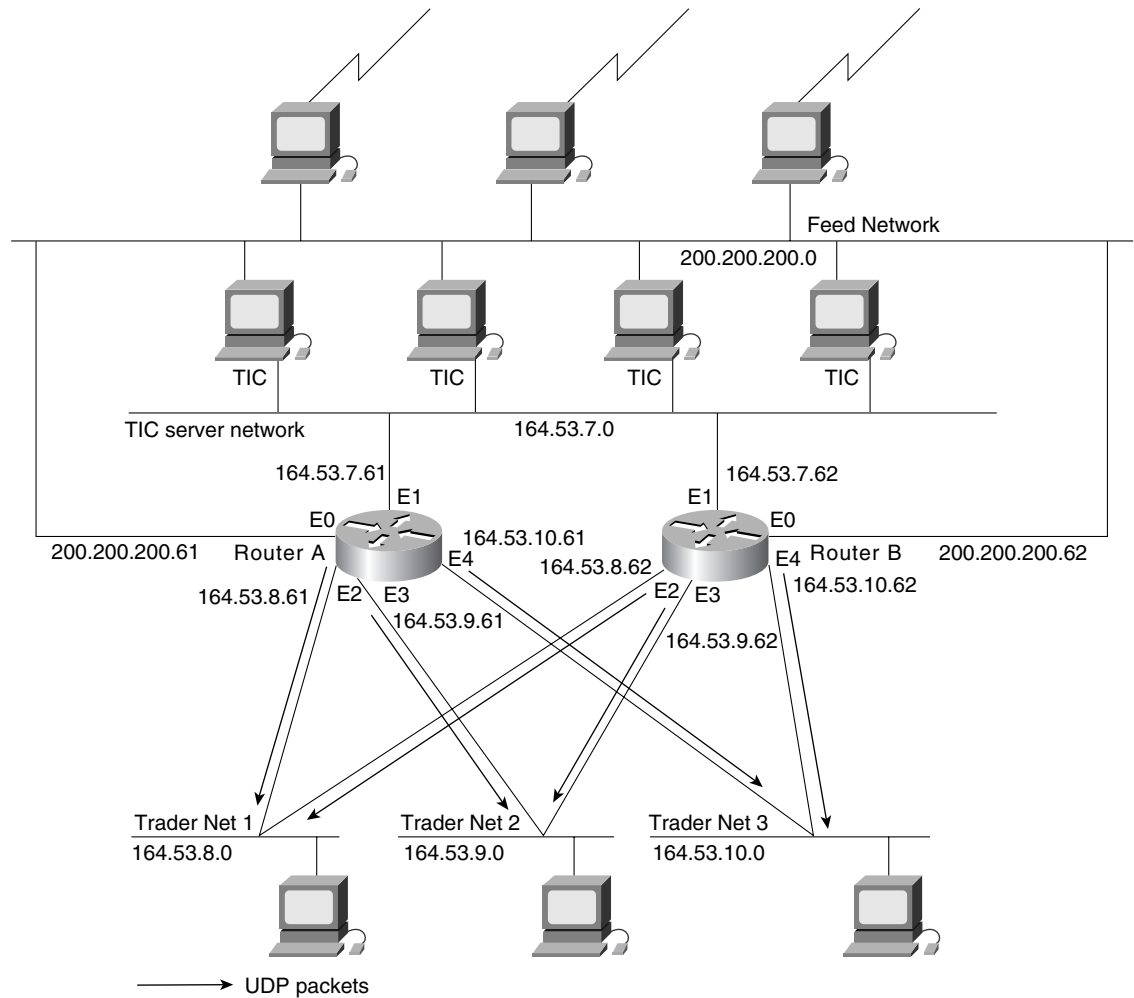
Note Regardless of whether you implement IP helper addressing or UDP flooding, you must use the **ip forward-protocol udp** global configuration command to enable the UDP forwarding. By default, the **ip forward-protocol udp** command enables forwarding for ports associated with the following protocols: Trivial File Transfer Protocol, Domain Name System, Time service, NetBIOS Name Server, NetBIOS Datagram Server, Boot Protocol, and Terminal Access Controller Access Control System. To enable forwarding for other ports, you must specify them as arguments to the **ip forward-protocol udp** command.

Implementing IP Helper Addressing

IP helper addressing is a form of static addressing that uses directed broadcasts to forward local and all-nets broadcasts to desired destinations within the internetwork.

To configure helper addressing, you must specify the **ip helper-address** command on every interface on every router that receives a broadcast that needs to be forwarded. On Router A and Router B, IP helper addresses can be configured to move data from the TIC server network to the trader networks. IP helper addressing is not the optimal solution for this type of topology because each router receives unnecessary broadcasts from the other router, as shown in Figure 6-2.

Figure 6-2 Flow of UDP packets from routers to trader networks using IP helper addressing.



In this case, Router A receives each broadcast sent by Router B *three times*, one for each segment, and Router B receives each broadcast sent by Router A three times, one for each segment. When each broadcast is received, the router must analyze it and determine that the broadcast does not need to be forwarded. As more segments are added to the network, the routers become overloaded with unnecessary traffic, which must be analyzed and discarded.

When IP helper addressing is used in this type of topology, no more than one router can be configured to forward UDP broadcasts (unless the receiving applications can handle duplicate broadcasts). This is because duplicate packets arrive on the trader network. This restriction limits redundancy in the design and can be undesirable in some implementations.

To send UDP broadcasts bidirectionally in this type of topology, a second **ip helper address** command must be applied to every router interface that receives UDP broadcasts. As more segments and devices are added to the network, more **ip helper address** commands are required to reach them, so the administration of these routers becomes more complex over time. Note, too, that bidirectional traffic in this topology significantly impacts router performance.

Although IP helper addressing is well-suited to nonredundant, nonparallel topologies that do not require a mechanism for controlling broadcast loops, in view of these drawbacks, IP helper addressing does not work well in this topology. To improve performance, network designers considered several other alternatives:

- *Setting the broadcast address on the TIC servers to all ones (255.255.255.255)*—This alternative was dismissed because the TIC servers have more than one interface, causing TIC broadcasts to be sent back onto the Feed network. In addition, some workstation implementations do not allow all ones broadcasts when multiple interfaces are present.
- *Setting the broadcast address of the TIC servers to the major net broadcast (164.53.0.0)*—This alternative was dismissed because the Sun TCP/IP implementation does not allow the use of major net broadcast addresses when the network is subnetted.
- *Eliminating the subnets and letting the workstations use Address Resolution Protocol (ARP) to learn addresses*—This alternative was dismissed because the TIC servers cannot quickly learn an alternative route in the event of a primary router failure.

With alternatives eliminated, the network designers turned to a simpler implementation that supports redundancy without duplicating packets and that ensures fast convergence and minimal loss of data when a router fails: UDP flooding.

Implementing UDP Flooding

UDP flooding uses the spanning tree algorithm to forward packets in a controlled manner. Bridging is enabled on each router interface for the sole purpose of building the spanning tree. The spanning tree prevents loops by stopping a broadcast from being forwarded out an interface on which the broadcast was received. The spanning tree also prevents packet duplication by placing certain interfaces in the blocked state (so that no packets are forwarded) and other interfaces in the forwarding state (so that packets that need to be forwarded are forwarded).

To enable UDP flooding, the router must be running software that supports transparent bridging and bridging must be configured on each interface that is to participate in the flooding. If bridging is not configured for an interface, the interface will receive broadcasts, but the router will not forward those broadcasts and will not use that interface as a destination for sending broadcasts received on a different interface.

Note Releases prior to Cisco Internetwork Operating System (Cisco IOS) Software Release 10.2 do not support flooding subnet broadcasts.

When configured for UDP flooding, the router uses the destination address specified by the **ip broadcast-address** command on the output interface to assign a destination address to a flooded UDP datagram. Thus, the destination address might change as the datagram propagates through the network. The source address, however, does not change.

With UDP flooding, both routers shown in Figure 19-1 use a spanning tree to control the network topology for the purpose of forwarding broadcasts. The key commands for enabling UDP flooding are as follows:

```
bridge group protocol protocol
ip forward-protocol spanning tree
bridge-group group input-type-list access-list-number
```

The **bridge protocol** command can specify either the **dec** keyword (for the DEC spanning-tree protocol) or the **ieee** keyword (for the IEEE Ethernet protocol). All routers in the network must enable the same spanning tree protocol. The **ip forward-protocol spanning tree** command uses the database created by the **bridge protocol** command. Only one broadcast packet arrives at each segment, and UDP broadcasts can traverse the network in both directions.

Note Because bridging is enabled only to build the spanning tree database, use access lists to prevent the spanning tree from forwarding non-UDP traffic. The configuration examples later in this chapter configure an access list that blocks all bridged packets.

To determine which interface forwards or blocks packets, the router configuration specifies a path cost for each interface. The default path cost for Ethernet is 100. Setting the path cost for each interface on Router B to 50 causes the spanning tree algorithm to place the interfaces in Router B in forwarding state. Given the higher path cost (100) for the interfaces in Router A, the interfaces in Router A are in the blocked state and do not forward the broadcasts. With these interface states, broadcast traffic flows through Router B. If Router B fails, the spanning tree algorithm will place the interfaces in Router A in the forwarding state, and Router A will forward broadcast traffic.

With one router forwarding broadcast traffic from the TIC server network to the trader networks, it is desirable to have the other forward unicast traffic. For that reason, each router enables the ICMP Router Discovery Protocol (IRDP), and each workstation on the trader networks runs the **irdp** daemon. On Router A, the **preference** keyword sets a higher IRDP preference than does the configuration for Router B, which causes each **irdp** daemon to use Router A as its preferred default gateway for unicast traffic forwarding. Users of those workstations can use **netstat -rn** to see how the routers are being used.

On the routers, the **holdtime**, **maxadvertinterval**, and **minadvertinterval** keywords reduce the advertising interval from the default so that the **irdp** daemons running on the hosts expect to see advertisements more frequently. With the advertising interval reduced, the workstations will adopt Router B more quickly if Router A becomes unavailable. With this configuration, when a router becomes unavailable, IRDP offers a convergence time of less than one minute.

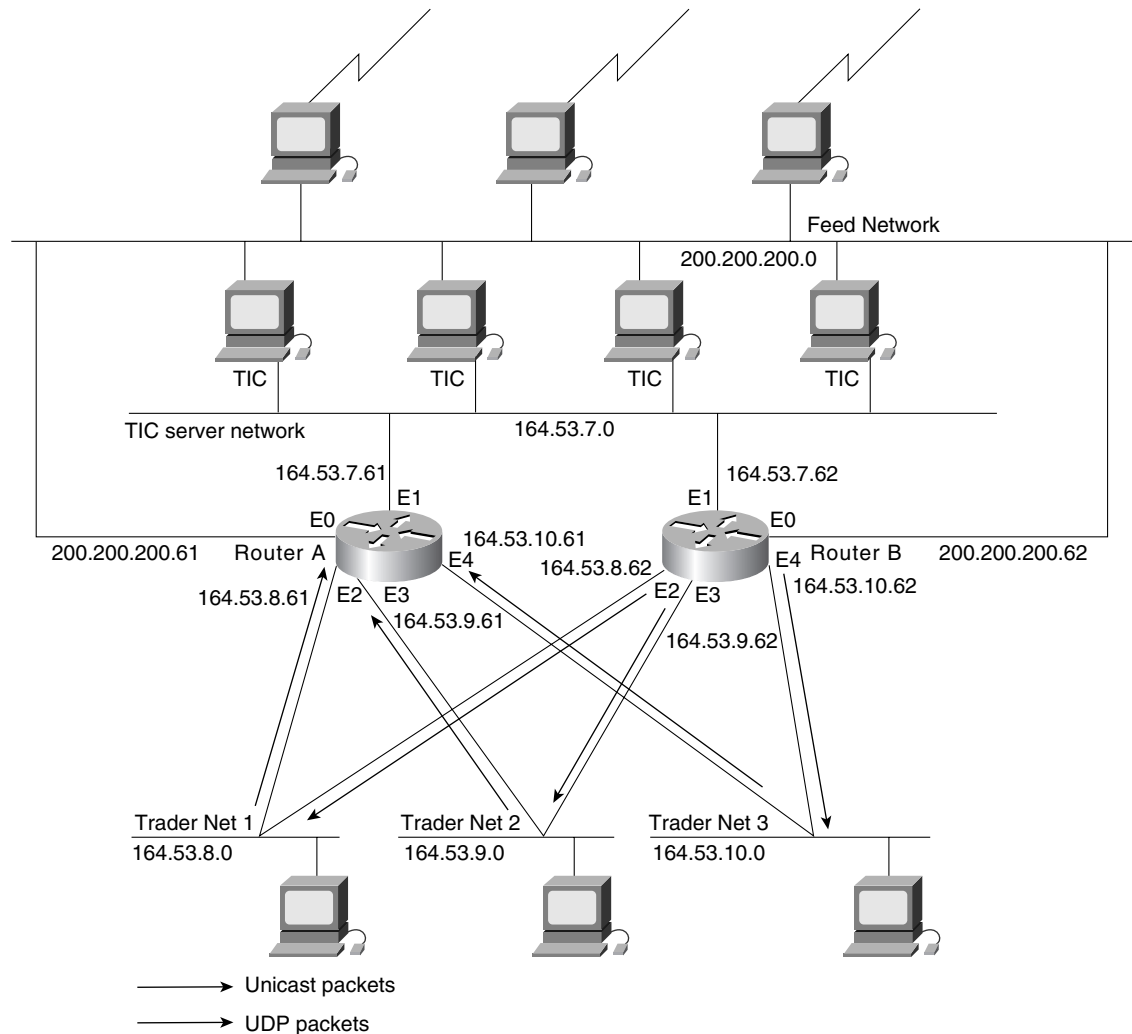
IRDP is preferred over the Routing Information Protocol (RIP) and default gateways for the following reasons:

- RIP takes longer to converge, typically from one to two minutes.
- Configuration of Router A as the default gateway on each Sun workstation on the trader networks would allow those Sun workstations to send unicast traffic to Router A, but would not provide an alternative route if Router A becomes unavailable.

Note Some workstation vendors include an **irdp** daemon with their operating systems. Source code for an **irdp** daemon is available by anonymous FTP at [ftp.cisco.com](ftp://ftp.cisco.com).

Figure 6-3 shows how data flows when the network is configured for UDP flooding.

Figure 6-3 Data flow with UDP flooding and IRDP.



Note This topology is broadcast intensive—broadcasts sometimes consume 20 percent of the Ethernet bandwidth. However, this is a favorable percentage when compared to the configuration of IP helper addressing, which, in the same network, causes broadcasts to consume up to 50 percent of the Ethernet bandwidth.

If the hosts on the trader networks do not support IRDP, the Hot Standby Routing Protocol (HSRP) can be used to select which router will handle unicast traffic. HSRP allows the standby router to take over quickly if the primary router becomes unavailable. For information about configuring HSRP, see Chapter 9, “Using HSRP for Fault-Tolerant IP Routing.”

By default, the router performs UDP flooding by process switching the UDP packets. To increase performance on AGS+ and Cisco 7000 series routers, enable fast switching of UDP packets by using the following command:

```
ip forward-protocol turbo-flood
```

Note Turbo flooding increases the amount of processing that is done at interrupt level, which increases the CPU load on the router. Turbo flooding may not be appropriate on routers that are already under high CPU load or that must also perform other CPU-intensive activities.

The following commands configure UDP flooding on Router A. Because this configuration does not specify a lower path cost than the default and because the configuration of Router B specifies a lower cost than the default with regard to UDP flooding, Router A acts as a backup to Router B. Because this configuration specifies an IRDP preference of 100 and because Router B specifies a IRDP preference of 90 (**ip irdp preference 90**), Router A forwards unicast traffic from the trader networks, and Router B is the backup for unicast traffic forwarding.

```
!Router A:
ip forward-protocol spanning-tree
ip forward-protocol udp 111
ip forward-protocol udp 3001
ip forward-protocol udp 3002
ip forward-protocol udp 3003
ip forward-protocol udp 3004
ip forward-protocol udp 3005
ip forward-protocol udp 3006
ip forward-protocol udp 5020
ip forward-protocol udp 5021
ip forward-protocol udp 5030
ip forward-protocol udp 5002
ip forward-protocol udp 1027
ip forward-protocol udp 657
!
interface ethernet 0
ip address 200.200.200.61 255.255.255.0
ip broadcast-address 200.200.200.255
no mop enabled
!
interface ethernet 1
ip address 164.53.7.61 255.255.255.192
ip broadcast-address 164.53.7.63
ip irdp
ip irdp maxadvertinterval 60
ip irdp minadvertinterval 45
ip irdp holdtime 60
ip irdp preference 100
bridge-group 1
bridge-group 1 input-type-list 201
no mop enabled
!
interface ethernet 2
ip address 164.53.8.61 255.255.255.192
ip broadcast-address 164.53.8.63
ip irdp
ip irdp maxadvertinterval 60
ip irdp minadvertinterval 45
ip irdp holdtime 60
ip irdp preference 100
bridge-group 1
bridge-group 1 input-type-list 201
no mop enabled
!
interface ethernet 3
ip address 164.53.9.61 255.255.255.192
ip broadcast-address 164.53.9.63
ip irdp
ip irdp maxadvertinterval 60
```

```

ip irdp minadvertinterval 45
ip irdp holdtime 60
ip irdp preference 100
bridge-group 1
bridge-group 1 input-type-list 201
no mop enabled
!
interface ethernet 4
ip address 164.53.10.61 255.255.255.192
ip broadcast-address 164.53.10.63
ip irdp
ip irdp maxadvertinterval 60
ip irdp minadvertinterval 45
ip irdp holdtime 60
ip irdp preference 100
bridge-group 1
bridge-group 1 input-type-list 201
no mop enabled
!
router igrp 1
network 164.53.0.0
!
ip name-server 255.255.255.255
snmp-server community public RW
snmp-server host 164.53.7.15 public
bridge 1 protocol dec
bridge 1 priority 255
access-list 201 deny 0xFFFF 0x0000

```

The following commands configure UDP flooding on Router B. Because this configuration specifies a lower path cost than the default (**bridge-group 1 path-cost 50**) and because the configuration of Router A accepts the default, Router B forwards UDP packets. Because this configuration specifies an IRDP preference of 90 (**ip irdp preference 90**) and because Router A specifies a IRDP preference of 100, Router B acts as the backup for Router A for forwarding unicast traffic from the trader networks.

```

!Router B
ip forward-protocol spanning-tree
ip forward-protocol udp 111
ip forward-protocol udp 3001
ip forward-protocol udp 3002
ip forward-protocol udp 3003
ip forward-protocol udp 3004
ip forward-protocol udp 3005
ip forward-protocol udp 3006
ip forward-protocol udp 5020
ip forward-protocol udp 5021
ip forward-protocol udp 5030
ip forward-protocol udp 5002
ip forward-protocol udp 1027
ip forward-protocol udp 657
!
interface ethernet 0
ip address 200.200.200.62 255.255.255.0
ip broadcast-address 200.200.200.255
no mop enabled
!
interface ethernet 1
ip address 164.53.7.62 255.255.255.192
ip broadcast-address 164.53.7.63
ip irdp
ip irdp maxadvertinterval 60
ip irdp minadvertinterval 45
ip irdp holdtime 60

```

```
ip irdp preference 90
bridge-group 1
bridge-group 1 path-cost 50
bridge-group 1 input-type-list 201
no mop enabled
!
interface ethernet 2
ip address 164.53.8.62 255.255.255.192
ip broadcast-address 164.53.8.63
ip irdp
ip irdp maxadvertinterval 60
ip irdp minadvertinterval 45
ip irdp holdtime 60
ip irdp preference 90
bridge-group 1
bridge-group 1 path-cost 50
bridge-group 1 input-type-list 201
no mop enabled
!
interface ethernet 3
ip address 164.53.9.62 255.255.255.192
ip broadcast-address 164.53.9.63
ip irdp
ip irdp maxadvertinterval 60
ip irdp minadvertinterval 45
ip irdp holdtime 60
ip irdp preference 90
bridge-group 1
bridge-group 1 path-cost 50
bridge-group 1 input-type-list 201
no mop enabled
!
interface ethernet 4
ip address 164.53.10.62 255.255.255.192
ip broadcast-address 164.53.10.63
ip irdp
ip irdp maxadvertinterval 60
ip irdp minadvertinterval 45
ip irdp holdtime 60
ip irdp preference 90
bridge-group 1
bridge-group 1 path-cost 50
bridge-group 1 input-type-list 201
no mop enabled
!
router igrp 1
network 164.53.0.0
!
ip name-server 255.255.255.255
snmp-server community public RW
snmp-server host 164.53.7.15 public
bridge 1 protocol dec
bridge 1 priority 255
access-list 201 deny 0xFFFF 0x0000
```

Note In releases prior to Cisco IOS Software Release 10.2, the spanning tree algorithm prevented the forwarding of local broadcast addresses, but allowed the forwarding of local secondary addresses. For that reason, when running a release prior to Cisco IOS Software Release 10.2, a secondary address must be specified for each Ethernet interface that will forward local broadcast address packets. The secondary address is used to forward packets, whereas the primary address is never used. In such configurations, the secondary addresses are assigned to an Interior Gateway Routing Protocol (IGRP) group instead of the primary address.

Summary

Although IP helper addressing is useful in networks that do not require redundancy, when configured in networks that feature redundancy, IP helper addressing results in packet duplication that severely reduces router and network performance.

By configuring UDP flooding, one router forwards UDP traffic without burdening the second router with duplicate packets. By dedicating one router to the task of forwarding UDP traffic, the second router becomes available for forwarding unicast traffic. At the same time, because each router is configured as the backup for the other router, redundancy is maintained; if either router fails, the other router can assume the work of the failed router without intervention from an operator. When compared with IP helper addressing, UDP flooding makes the most efficient use of router resources.

