

Using HSRP for Fault-Tolerant IP Routing


This case study examines Cisco's Hot Standby Routing Protocol (HSRP), which provides automatic router backup when you configure it on Cisco routers that run the Internet Protocol (IP) over Ethernet, Fiber Distributed Data Interface (FDDI), and Token Ring local-area networks (LANs). HSRP is compatible with Novell's Internetwork Packet Exchange (IPX), AppleTalk, and Banyan VINES, and it is compatible with DECnet and Xerox Network Systems (XNS) in certain configurations.

Note Banyan VINES serverless clients do not respond well to topology changes (regardless of whether HSRP is configured). This case study describes the effect of topology changes in networks that include Banyan VINES serverless clients.

For IP, HSRP allows one router to automatically assume the function of the second router if the second router fails. HSRP is particularly useful when the users on one subnet require continuous access to resources in the network.

Consider the network shown in Figure 9-1. Router A is responsible for handling packets between the Tokyo segment and the Paris segment, and Router B is responsible for handling packets between the Tokyo segment and the New York segment. If the connection between Routers A and C goes down or if either router becomes unavailable, fast converging routing protocols, such as the Enhanced Interior Gateway Routing Protocol (Enhanced IGRP) and Open Shortest Path First (OSPF) can respond within seconds so that Router B is prepared to transfer packets that would otherwise have gone through Router A.

Figure 9-1 Typical WAN



However, in spite of fast convergence, if the connection between Router A and Router C goes down, or if either router becomes unavailable, the user Pat on the Tokyo segment might not be able to communicate with the user Marceau even after the routing protocol has converged. That's because IP hosts, such as Pat's workstation, usually do not participate in routing protocols. Instead, they are configured statically with the address of a single router, such as Router A. Until someone manually modifies the configuration of Pat's host to use the address of Router B instead of Router A, Pat cannot communicate with Marceau.


Some IP hosts use proxy Address Resolution Protocol (ARP) to select a router. If Pat's workstation were running proxy ARP, it would send an ARP request for the IP address of Marceau's workstation. Router A would reply on behalf of Marceau's workstation and would give to Pat's workstation its own media access control (MAC) address (instead of the IP address of Marceau's workstation). With proxy ARP, Pat's workstation behaves as if Marceau's workstation were connected to the same segment of the network as Pat's workstation. If Router A fails, Pat's workstation will continue to send packets destined for Marceau's workstation to the MAC address of Router A even though those packets have nowhere to go and are lost. Pat either waits for ARP to acquire the MAC address of Router B by sending another ARP request or reboots the workstation to force it to send an ARP request. In either case, for a significant period of time, Pat cannot communicate with Marceau—even though the routing protocol has converged, and Router B is prepared to transfer packets that would otherwise go through Router A.

Some IP hosts use the Routing Information Protocol (RIP) to discover routers. The drawback of using RIP is that it is slow to adapt to changes in the topology. If Pat's workstation is configured to use RIP, 3 to 10 minutes might elapse before RIP makes another router available.

Some newer IP hosts use the ICMP Router Discovery Protocol (IRDP) to find a new router when a route becomes unavailable. A host that runs IRDP listens for "hello" multicast messages from its configured router and uses an alternate router when it no longer receives those hello messages. If Pat's workstation were running IRDP, it would detect that Router A is no longer sending hello messages and would start sending its packets to Router B.

For IP hosts that do not support IRDP, Cisco's HSRP provides a way to keep communicating when a router becomes unavailable. HSRP allows two or more HSRP-configured routers to use the MAC address and IP network address of a virtual router. The virtual router does not physically exist—instead, it represents the common target for routers that are configured to provide backup to each other. Figure 9-2 shows the Tokyo segment of the WAN as it might be configured for HSRP. Each actual router is configured with the MAC address and the IP network address of the virtual router.

Figure 9-2 HSRP Addressing on the Tokyo Segment



In Figure 9-2, the MAC address of the virtual router is 0000.0c07.ac01. When you configure HSRP, the router automatically selects one of the virtual MAC addresses from a range of addresses in the Cisco IOS software that is within the range of Cisco's MAC address block. Ethernet and FDDI LANs use one of the preassigned MAC addresses as a virtual MAC address. Token Ring LANs use a functional address as a virtual MAC address.

In Figure 9-2, instead of configuring the hosts on network 192.1.1.0 with the IP address of Router A, they are configured with the IP address of the virtual router as their default router. When Pat's workstation sends packets to Marceau's workstation on the Paris segment, it sends them to the MAC address of the virtual router.

In Figure 9-2, Router A is configured as the active router. It is configured with the IP address and MAC address of the virtual router and sends any packets addressed to the virtual router out interface 1 to the Paris segment. As the standby router, Router B is also configured with the IP address and MAC address of the virtual router. If for any reason Router A stops transferring packets, the routing protocol converges, and Router B assumes the duties of Router A and becomes the active router. That is, Router B now responds to the virtual IP address and the virtual MAC address. Pat's workstation continues to use the IP address of the virtual router to address packets destined for Marceau's workstation, which Router B receives and sends to the Paris segment via the New York segment. Until Router A resumes operation, HSRP allows Router B to provide uninterrupted service to the users on the Tokyo segment that need to communicate with users on the Paris segment. While it is the active router, Router B continues to perform its normal function: handling packets between the Tokyo segment and the New York segment.

HSRP also works when the hosts are configured for proxy ARP. When the active HSRP router receives an ARP request for a host that is not on the local LAN, the router replies with the MAC address of the virtual router. If the active router becomes unavailable or its connection to the remote LAN goes down, the router that becomes the active router receives packets addressed to the virtual router and transfers them accordingly.

Note You can configure HSRP on any Cisco router that is running Cisco Internetwork Operating System (Cisco IOS) Software Release 10.0 or later. If you configure HSRP for one Cisco router on a Token Ring LAN, all Cisco routers on that LAN must run Cisco IOS Software Release 10.0 or later. Cisco IOS Software Releases 10.2(9), 10.3(6), and 11.0(2) allow standby IP addresses to respond to ping requests. Cisco Software Release 11.0(3)(1) provides improved support for the use of secondary IP addresses with HSRP.

How HSRP Works

HSRP uses a priority scheme to determine which HSRP-configured router is to be the default active router. To configure a router as the active router, you assign it a priority that is higher than the priority of all the other HSRP-configured routers. The default priority is 100, so if you configure just one router to have a higher priority, that router will be the default active router.

HSRP works by the exchange of multicast messages that advertise priority among HSRP-configured routers. When the active router fails to send a hello message within a configurable period of time, the standby router with the highest priority becomes the active router. The transition of packet-forwarding functions between routers is completely transparent to all hosts on the network.

HSRP-configured routers exchange three types of multicast messages:

- Hello—The hello message conveys to other HSRP routers the router's HSRP priority and state information. By default, an HSRP router sends hello messages every three seconds.
- Coup—When a standby router assumes the function of the active router, it sends a coup message.
- Resign. A router that is the active router sends this message when it is about to shut down or when a router that has a higher priority sends a hello message.

At any one time, HSRP-configured routers are in one of the following states:

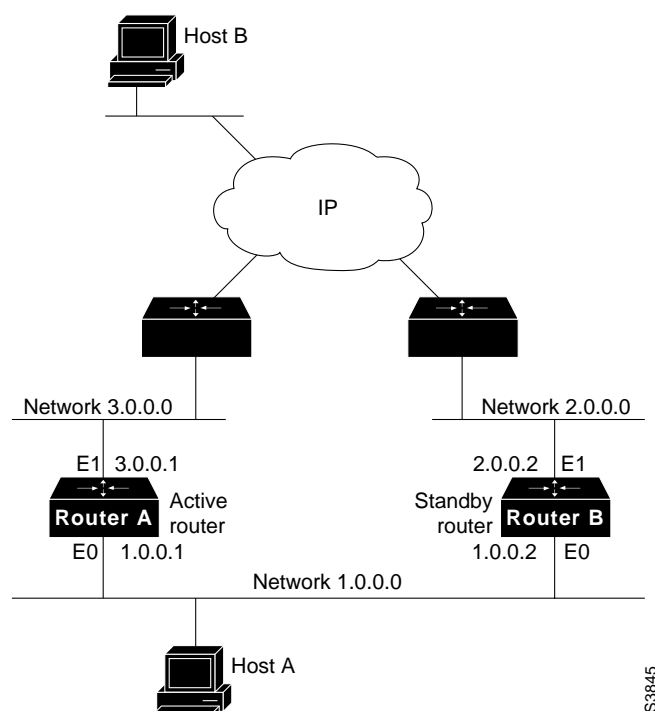
- Active—The router is performing packet-transfer functions.
- Standby—The router is prepared to assume packet-transfer functions if the active router fails.
- Speaking and listening—The router is sending and receiving hello messages.
- Listening—The router is receiving hello messages.

Note When configured on AGS, AGS+, and Cisco 7000 series routers, HSRP takes advantage of special hardware features that are not available on other Cisco routers. This means that HSRP operates in a slightly different way on these routers. For an example, see the “HSRP and Routed Protocols” section later in this chapter.

Configuring HSRP

Figure 9-3 shows the topology of a IP network in which two routers are configured for HSRP.

Figure 9-3 Example of a Network Configured for HSRP



S3845

All hosts on the network are configured to use the IP address of the virtual router (in this case, 1.0.0.3) as the default gateway. The command for configuring the default gateway depends on the host's operating system, TCP/IP implementation, and configuration.

Note The configurations shown in this case study use the Enhanced IGRP routing protocol. HSRP can be used with any routing protocol supported by the Cisco IOS software. Some configurations that use HSRP still require a routing protocol to converge when a topology change occurs. The standby router becomes active, but connectivity does not occur until the protocol converges.

Here is the configuration for Router A:

```
hostname RouterA
!
interface ethernet 0
ip address 1.0.0.1 255.0.0.0
standby 1 ip 1.0.0.3
standby 1 preempt
standby 1 priority 110
standby 1 authentication denmark
standby 1 timers 5 15
!
interface ethernet 1
ip address 3.0.0.1 255.0.0.0
!
router eigrp 1
network 1.0.0.0
network 3.0.0.0
```

Here is the configuration for Router B:

```
hostname RouterB
!
interface ethernet 0
ip address 1.0.0.2 255.0.0.0
standby 1 ip 1.0.0.3
standby 1 preempt
standby 1 authentication denmark
standby 1 timers 5 15
!
interface ethernet 1
ip address 2.0.0.2 255.0.0.0
!
router eigrp 1
network 1.0.0.0
network 2.0.0.0
```

The **standby ip** interface configuration command enables HSRP and establishes 1.0.0.3 as the IP address of the virtual router. The configurations of both routers include this command so that both routers share the same virtual IP address. The 1 establishes Hot Standby group 1. (If you do not specify a group number, the default is group 0.) The configuration for at least one of the routers in the Hot Standby group must specify the IP address of the virtual router; specifying the IP address of the virtual router is optional for other routers in the same Hot Standby group.

The **standby preempt** interface configuration command allows the router to become the active router when its priority is higher than all other HSRP-configured routers in this Hot Standby group. The configurations of both routers include this command so that each router can be the standby

router for the other router. The 1 indicates that this command applies to Hot Standby group 1. If you do not use the **standby preempt** command in the configuration for a router, that router cannot become the active router.

The **standby priority** interface configuration command sets the router's HSRP priority to 110, which is higher than the default priority of 100. Only the configuration of Router A includes this command, which makes Router A the default active router. The 1 indicates that this command applies to Hot Standby group 1.


The **standby authentication** interface configuration command establishes an authentication string whose value is an unencrypted eight-character string that is incorporated in each HSRP multicast message. This command is optional. If you chose to use it, each HSRP-configured router in the group should use the same string so that each router can authenticate the source of the HSRP messages that it receives. The "1" indicates that this command applies to Hot Standby group 1.

The **standby timers** interface configuration command sets the interval in seconds between hello messages (called the *hello time*) to 5 seconds and sets the duration in seconds that a router waits before it declares the active router to be down (called the *hold time*) to 8 seconds. (The defaults are 3 and 10 seconds, respectively.) If you decide to modify the default values, you must configure each router to use the same hello time and hold time. The 1 indicates that this command applies to Hot Standby group 1.

Note There can be up to 255 Hot Standby groups on any Ethernet or FDDI LAN. There can be no more than three Hot Standby groups on any Token Ring LAN.

Configuring Multiple Hot Standby Groups

Multigroup HSRP (MHSRP) is an extension of HSRP that allows a single router interface to belong to more than one Hot Standby group. MHSRP requires the use of Cisco IOS Software Release 10.3 or later and is supported only on routers that have special hardware that allows them to associate an Ethernet interface with multiple unicast Media Access Control (MAC) addresses. These routers are the AGS and AGS+ routers and any router in the Cisco 7000 series. The special hardware allows you to configure a single interface in an AGS, AGS+, or Cisco 7000 series router so that the router is the backup router for more than one Hot Standby group, as shown in Figure 9-4.

Figure 9-4 Example of Hot Standby Groups

In Figure 9-4, the Ethernet interface 0 of Router A belongs to group 1. Ethernet interface 0 of Router B belongs to groups 1, 2, and 3. The Ethernet interface 0 of Router C belongs to group 2, and the Ethernet interface 0 of Router D belongs to group 3. When you establish groups, you might want to align them along departmental organizations. In this case, group 1 might support the Engineering Department, group 2 might support the Manufacturing Department, and group 3 might support the Finance Department.

Router B is configured as the active router for groups 1 and 2 and as the standby router for group 3. Router D is configured as the active router for group 3. If Router D fails for any reason, Router B will assume the packet-transfer functions of Router D and will maintain the ability of users in the Finance Department to access data on other subnets.

Here is the configuration for Router A:

```
hostname RouterA
!
interface ethernet 0
ip address 1.0.0.1 255.0.0.0
standby 1 ip 1.0.0.5
standby authentication sclara
!
interface serial 0
ip address 2.0.0.1 255.0.0.0
!
router eigrp 1
network 1.0.0.0
network 2.0.0.0
```

Here is the configuration for Router B, which must be an AGS, AGS+, or Cisco 7000 series router:

```
hostname RouterB
!
interface ethernet 0
ip address 1.0.0.2 255.0.0.0
standby 1 ip 1.0.0.5
standby 1 priority 110
standby 1 preempt
standby 1 authentication sclara
```

```
standby 2 ip 1.0.0.6
standby 2 priority 110
standby 2 preempt
standby 2 authentication mtview
standby 3 ip 1.0.0.7
standby 3 preempt
standby 3 authentication svale
!
interface serial 0
ip address 3.0.0.1 255.0.0.0
!
router eigrp 1
network 1.0.0.0
network 3.0.0.0
```

Here is the configuration for Router C:


```
hostname RouterC
!
interface ethernet 0
ip address 1.0.0.3 255.0.0.0
standby 2 ip 1.0.0.6
standby 2 authentication mtview
!
interface serial 0
ip address 4.0.0.1 255.0.0.0
!
router eigrp 1
network 1.0.0.0
network 4.0.0.0
```

Here is the configuration for Router D:

```
hostname RouterD
!
interface ethernet 0
ip address 1.0.0.4 255.0.0.0
standby 3 ip 1.0.0.7
standby 1 priority 110
standby 1 preempt
standby 3 authentication svale
!
interface serial 0
ip address 4.0.0.1 255.0.0.0
!
router eigrp 1
network 1.0.0.0
network 5.0.0.0
```

Interface Tracking

For both HSRP and MHSRP, you can use the tracking feature to adjust the Hot Standby priority of a router based on whether certain of the router's interfaces are available. When a tracked interface becomes unavailable, the HSRP priority of the router is decreased. You can use tracking to automatically reduce the likelihood that a router that already has an unavailable key interface will become the active router. To configure tracking, use the **standby track** interface configuration command. Figure 9-5 shows a network for which tracking is configured.

Figure 9-5 Tracking Example

In Figure 9-5, Router A is configured as the active router. Routers B and C are configured as standby routers for Router A.

Here is the configuration for Router A:

```
hostname RouterA
!
interface ethernet 0
ip address 1.0.0.1 255.0.0.0
standby 1 ip 1.0.0.4
standby 1 preempt
standby 1 priority 110
standby authentication microdot
!
interface serial 0
ip address 2.0.0.1 255.0.0.0
!
router eigrp 1
network 1.0.0.0
network 3.0.0.0
```

The **standby ip** interface configuration command enables HSRP and establishes 1.0.0.4 as the IP address of the virtual router. The 1 establishes Hot Standby group 1.

The **standby preempt** interface configuration command allows Router A to become the active router when its priority is higher than all other HSRP-configured routers in the Hot Standby group.

The **standby priority** interface configuration command sets the router's HSRP priority to 110, which is highest priority assigned to the three routers in this example. Because Router A has the highest priority, it is the active router under normal operation.

Here is the configuration for Router B:

```
hostname RouterB
!
interface ethernet 0
ip address 1.0.0.2 255.0.0.0
standby 1 ip 1.0.0.4
standby 1 preempt
standby 1 priority 105
standby track serial 0
standby 1 authentication microdot
```

```
interface serial 0
ip address 3.0.0.1 255.0.0.0
!
router eigrp 1
network 1.0.0.0
network 2.0.0.0
```

The **standby preempt** interface configuration command allows Router B to become the active router immediately if its priority is highest, even before the current active router fails. The **standby priority** interface configuration command specifies a priority of 105 (lower than the priority of Router A and higher than the priority of Router C), so Router B is a standby router.

The **standby track** interface configuration command causes Ethernet interface 0 to track serial interface 0. If serial interface 0 becomes unavailable, the priority of Router B is reduced by 10 (the default).

Here is the configuration for Router C:

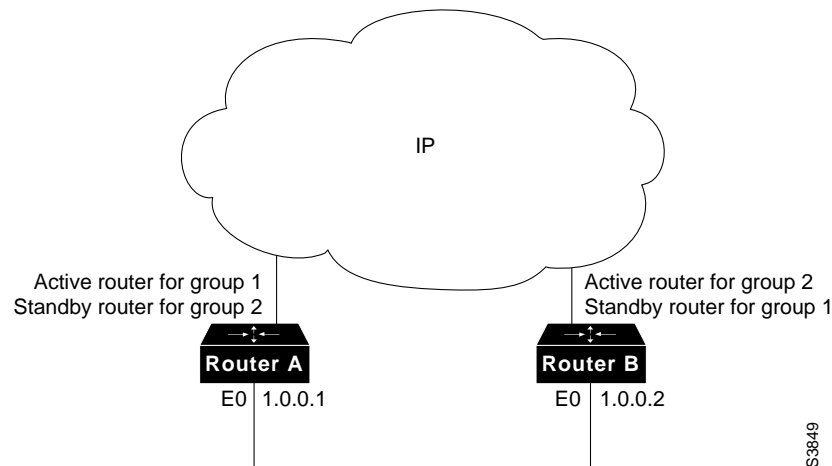
```
hostname RouterC
!
interface ethernet 0
ip address 1.0.0.3 255.0 0.0
standby 1 ip 1.0.0.4
standby 1 preempt
standby 1 priority
standby track serial 0
standby 1 authentication microdot
!
interface serial 0
ip address 4.0.0.1 255.0.0.0
!
router eigrp 1
network 1.0.0.0
network 4.0.0.0
```

The **standby preempt** interface configuration command allows Router C to become the active router if its priority is highest when the active router fails. The **standby priority** interface configuration command does not specify a priority, so its priority is 100 (the default).

If Router A becomes unavailable, and if serial interface 0 on Router B is available, Router B (with its priority of 105) will become the active router. However, if serial interface 0 on Router B becomes unavailable before Router A becomes unavailable, the HSRP priority of Router B will be reduced from 105 to 95. If Router A then becomes unavailable, Router C (whose priority is 100) will become the active router.

Load Sharing

You can use HSRP or MHSRP when you configure load sharing. In Figure 9-6, half of the workstations on the LAN are configured for Router A, and half of the workstations are configured for Router B.

Figure 9-6 Load Sharing Example

Here is a partial configuration for Router A:

```
hostname RouterA
!
interface ethernet 0
ip address 1.0.0.1 255.0.0.0
standby 1 ip 1.0.0.3
standby 1 priority 110
standby 1 preempt
standby 2 ip 1.0.0.4
standby 2 preempt
```

Here is a partial configuration for Router B:

```
hostname RouterB
!
interface ethernet 0
ip address 1.0.0.2 255.0.0.0
standby 1 ip 1.0.0.3
standby 1 preempt
standby 2 ip 1.0.0.4
standby 2 priority 110
standby 2 preempt
```

Together, the configuration files for Routers A and B establish two Hot Standby groups. For group 1, Router A is the default active router, and Router B is the standby router. For group 2, Router B is the default active router, and Router A is the standby router. During normal operation, the two routers share the IP traffic load. When either router becomes unavailable, the other router transitions to active and assumes the packet-transfer functions of the router that is unavailable. The **standby preempt** interface configuration commands are necessary so that if a router goes down and then comes back up, preemption occurs and restores load sharing.

HSRP and Routed Protocols

This section describes the interaction between HSRP and the following routed protocols:

- AppleTalk, Banyan VINES, and Novell IPX
- DECnet and XNS

AppleTalk, Banyan VINES, and Novell IPX

You can configure HSRP in networks that, in addition to IP, run AppleTalk, Banyan VINES, and Novell IPX. AppleTalk and Novell IPX continue to function when the standby router becomes the active router, but they take time to adapt to topology changes. In general, AppleTalk hosts discover a new active router in less than 30 seconds. Novell 4.x hosts discover a new active router in 10 seconds, on average. Novell 2.x or Novell 3.x hosts might require more time to adapt.

Regardless of whether HSRP is configured, Banyan VINES does not respond well to topology changes. When HSRP is configured, the effect of a topology change varies, depending on the type of router that becomes the active router.

When the active router becomes unavailable, or its connection to the network goes down, all Banyan VINES sessions that rely on that router stop and must be reinitiated. If an AGS, AGS+, or Cisco 7000 series router becomes the active router, Banyan VINES traffic flowing through that router is not affected as it transitions from standby to active. That is because these routers have special hardware that allows them to have more than one MAC address at the same time. If the router that becomes the active router is *not* an AGS, AGS+, or Cisco 7000 series router, Banyan VINES traffic flowing through that router pauses and resumes after no more than 90 seconds while the router transitions from standby to active.

Regardless of which type of router becomes the active router, any Banyan VINES serverless clients that obtained their network-layer address from the unavailable router might need to reboot in order to obtain another network-layer address.

DECnet and XNS

DECnet and XNS are compatible with HSRP and MHSRP over Ethernet, FDDI, and Token Ring on the Cisco 7000 and Cisco 7500 routers. Some constraints apply when HSRP and MHSRP are configured on other routers, such as the Cisco 2500, Cisco 3000, Cisco 4000, and Cisco 4500 series routers, which do not have the hardware required to support multiple MAC addresses. Table 9-1 identifies the supported and unsupported combinations.

Table 9-1 HSRP and MSHRP Compatability with DECnet and XNS

Protocol Combination Per Interface	Cisco2 500	Cisco 3000	Cisco4 000	Cisco4 500	Cisco7 000	Cisco7 500
MHSRP with or without DECnet or XNS	No	No	No	No	Yes	Yes
HSRP without DECnet or XNS	Yes	Yes	Yes	Yes	Yes	Yes
HSRP with DECnet or XNS	No	No	No	No	Yes	Yes

Summary

HSRP and MHSRP provide fault tolerant routing of IP packets for networks that require nonstop access by hosts on all segments to resources on all segments. To provide fault tolerance, HSRP and MHSRP require a routing protocol that converges rapidly, such as Enhanced Interior Gateway Routing Protocol (Enhanced IGRP). A fast-converging protocol ensures that router state changes propagate fast enough to make the transition from standby to active mode transparent to network users.