

HSRP Support

Features, Usage, and Caveats Detailed

Background

The Need for HSRP

Because of the way that IP hosts determine where to forward their packets, there needs to be a way to enable quick and efficient redundancy or fail-over. Without HSRP the failure of a single gateway could isolate hosts. The ways in which IP hosts can determine where to forward packets are listed below:

1. The configuration of a default gateway. This is the simplest and most often used method. Using this method a host is statically configured to know the IP address of its default router. However, if that router should become unavailable, the host will no longer be able to communicate with devices off of the local LAN segment even if there is another router available.
2. The use of Proxy ARP. A host can discover a router by using ARP to find the MAC address of any hosts which are not on its directly connected LAN. A router which is configured to support Proxy ARP will answer ARP requests with its own MAC address when it has a specific route for these addresses. Unfortunately, many hosts have long time out values (or no timeout values at all) on their ARP caches. As a result, if a router should become unavailable, the host will continue to attempt to send traffic for these hosts to the router which originally sent the proxy ARP reply.
3. ICMP Router Discover Protocol (IRDP). IP hosts can use the Router Discovery Protocol to listen to router hellos. This allows a host to quickly adapt to changes in network topology. However, only a very small number of hosts have implementations of IRDP.
4. RIP. Some IP hosts use RIP to discover routers. These hosts will adapt to topology changes as RIP converges. Very few hosts fall into the category of devices which run RIP, and very few network administrators would like to deal with the overhead of another routing protocol in their networks or the hassle of having all of their hosts relying on a routing protocol for connectivity. Even if they will take this hit, the convergence of RIP can take over a minute.

Since none of these mechanisms are satisfactory in the majority of networking situations, Cisco has developed the Hot Standby Router Protocol to allow hosts to adapt to network topology changes almost immediately without requiring hosts to run any special software. HSRP is used in conjunction with the configuration of a default gateway in the host devices. This makes the protocol easy to use in any networking environment and provides redundancy for the critical first hop, or gateway.

General Functionality of HSRP

The Cisco innovation Hot Standby Router Protocol (HSRP) was introduced in IOS 10.0. HSRP brings fast re-routing technology to the desktop. HSRP enables a set of routers to work together to present the appearance of a single virtual router or default gateway to the hosts on a LAN. HSRP is particularly useful in environments where critical applications are running and fault-tolerant

networks have been designed. By sharing an IP address and a MAC address two or more routers acting as one virtual router are able to seamlessly assume the routing responsibility in the case of a defined event or the unexpected failure. This enables hosts on a LAN to continue to forward IP packets to a consistent IP and MAC address enabling the changeover of devices doing the routing to be transparent to them and their sessions.

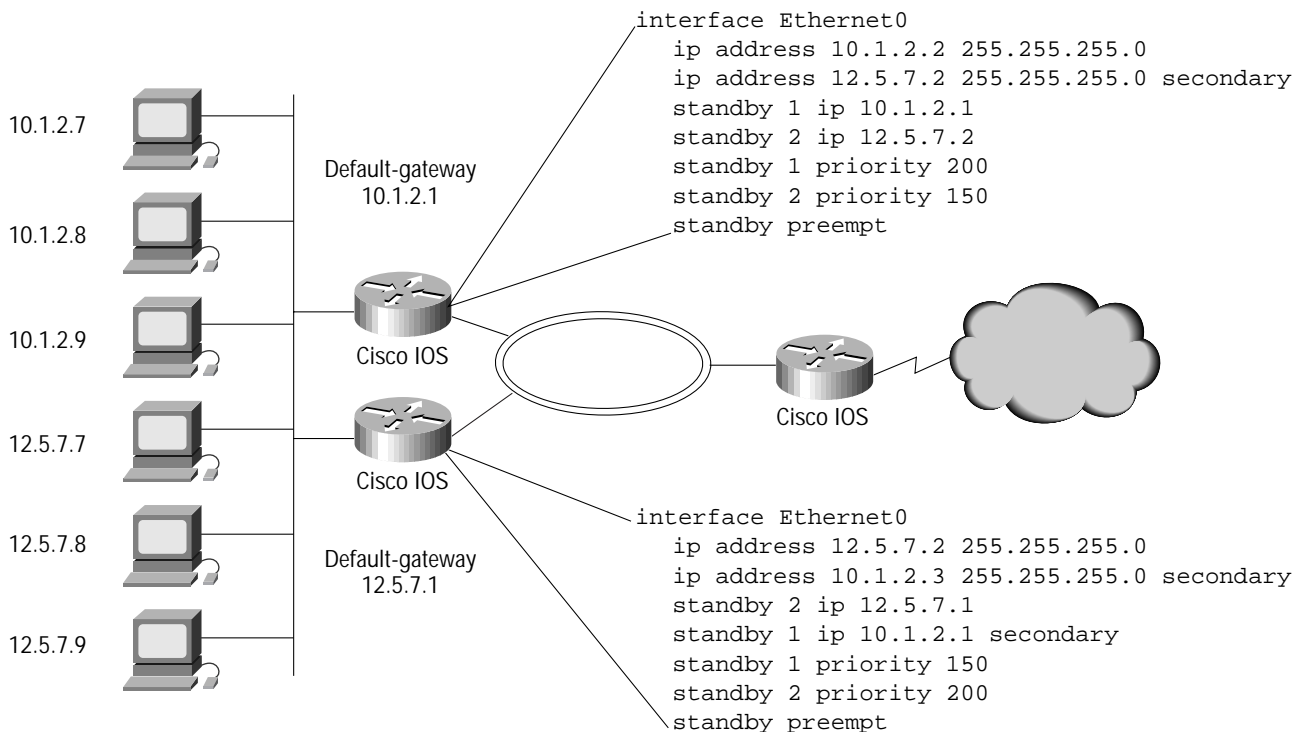
HSRP works by allowing an administrator to configure Hot Standby Groups to share responsibility for an IP address. Each router can be given a priority to enable an administrator to weight the prioritization of routers for active router selection. One of the routers in each group will be elected to be the active forwarder and one the stand-by router to stand ready to take over this functionality. This is done according to the router's configured priorities. The router with the highest priority wins and, in the case of a tie in priority, the greater value of their configured IP addresses will break the tie. Other routers in this group will monitor the active and stand-by routers' status to enable further fault tolerance. All HSRP routers participating in a standby group will watch for hello packets from the active and the standby routers. From the active router in the group they will all learn the hello and dead timer as well as the standby IP address to be shared, if these parameters are not explicitly configured on each individual router. If the active router becomes unavailable due to scheduled maintenance, power failure, or other reasons the stand-by can assume this functionality transparently within a few seconds. This will occur if the dead timer is reached, by missing three successive hello packets, and the standby router will promptly take over the virtual addresses, identity and responsibility.

Added Features

Multiple HSRP Groups & Secondary Addresses

A later enhancement which was added in IOS 10.3 allows Multiple HSRP (M-HSRP) groups to be configured per interface. This feature further enables redundancy and load-sharing within networks. This allows redundant routers to be more fully utilized. While a router is actively forwarding traffic for one HSRP group, it can be in standby or listen for another group enabling load-sharing of LAN originated traffic rather than having a standby router not be utilized. The use of HSRP for secondary addresses is supported beginning in IOS 10.3 also, further expanding the usability of HSRP for real networks. The following is an example of where both of these features would be beneficial.

Figure 1

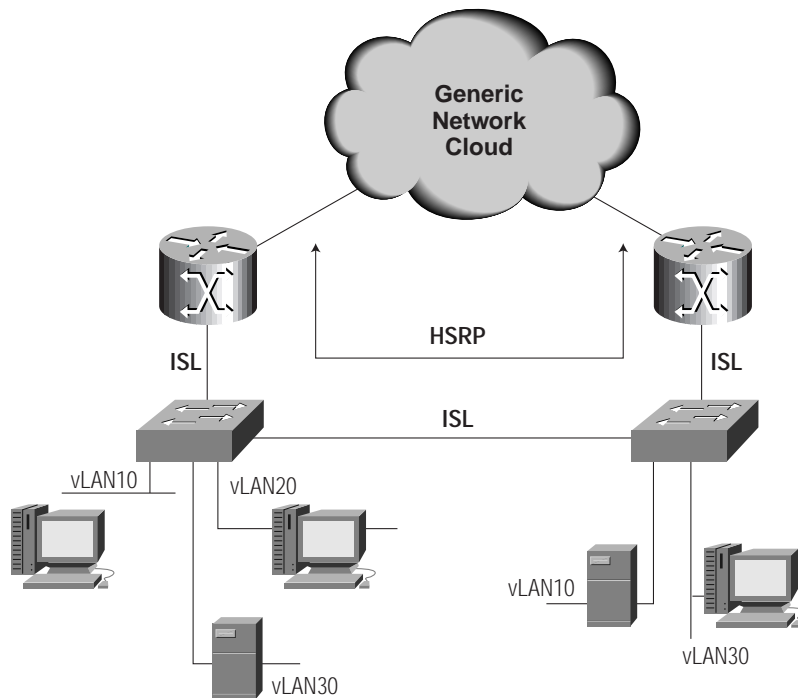


In this case each of the router on the LAN can be forwarding traffic for one group and be in standby for another, even though the IP addresses of the two subnets are different. This enables the backup of each of the active routers while enabling load sharing while both are available.

Increased Media and Interface Support

Since 10.0 the HSRP functionality has been available on Ethernet, Token Ring and FDDI. When Fast Ethernet and ATM interfaces became available those interface types were also supported by HSRP. This will be followed by gigabit Ethernet as well. Virtual LANS (vLANs) allow logical network topologies to overlay the physical switched infrastructure such that any arbitrary collection of LAN ports can be combined into an autonomous user group or community of interest. In 11.1 v LAN support was added for IEEE 802.10 SDE. Starting with IOS 11.3, HSRP functionality is also available on other vLANs, including Ethernet configured for Cisco's Inter Switch Link (ISL), or FastEtherchannel LANE encapsulations for Ethernet and Token Ring are also supported. HSRP will be supported in conjunction with IEEE 802.1Q vLANs also.

Figure 2



The use of vLAN encapsulations in conjunction with HSRP requires higher layers of software to be aware of MAC addresses embedded in the vLAN framing. Because of this IOS needed to be enhanced to deal with the MAC logic of HSRP in higher layers of software. Therefore, there are no issues with the use of the same MAC address on each several of the vLAN interfaces or on different interfaces in conjunction with the RSM or any other platform, since they are on logically separate interfaces with respect to the software. However, with the use of vLANs, there can exist network typologies, where the HSRP peers are separated by different media and devices, such that the takeover of routing responsibility by a standby router might not be immediately visible to all host devices. This can be the case because of different aging timers and cached entries in some switching devices and the existence of devices between the peers.

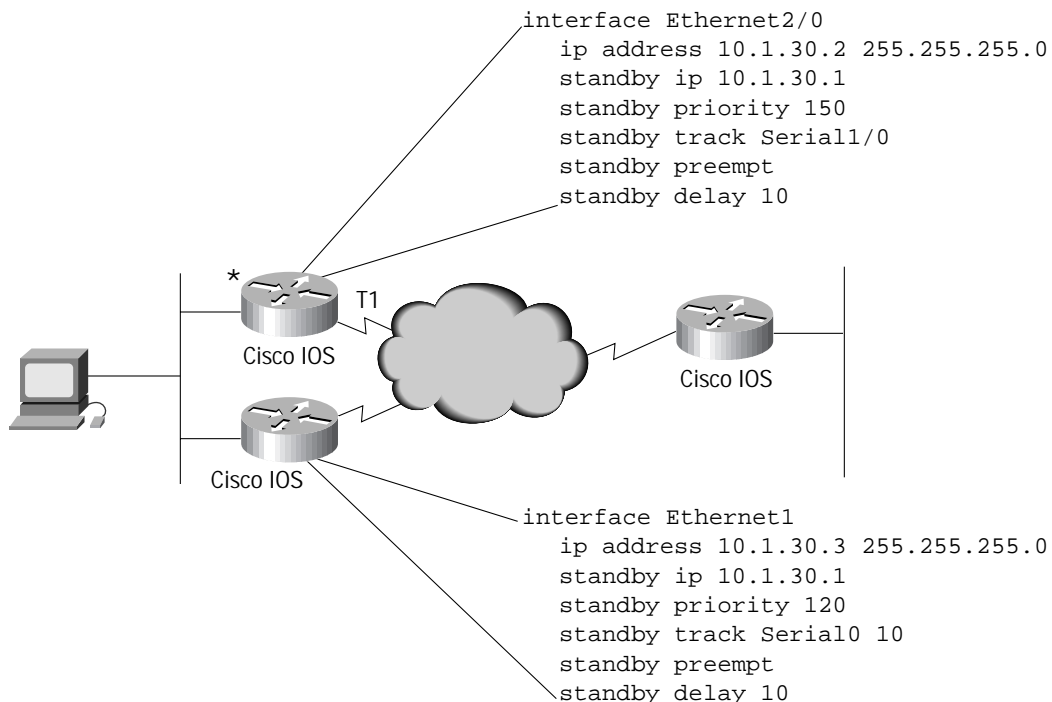
Preempt, the Tracking Feature and Preempt Delay

The preempt feature enables a router with the highest priority to immediately assume the active role at any time. Note that this does not apply to the higher IP address values used to break a tie when two routers originally vie for the active router role when they boot-up. Only a higher priority will allow a router to utilize the preempt mechanism and become the active router. Otherwise the default behavior of HSRP is for the standby HSRP router to wait until it does not receive three hello successive hello packets from the active HSRP router before it takes over active functionality, regardless of their respective priorities.

The tracking feature, added in IOS 10.3, allows administrators to specify another interface on the router for the HSRP process to monitor in order to alter the HSRP priority for a given group. If a specified interface's line protocol goes down the HSRP priority can be reduced. This means that another suitable HSRP router, with a higher priority, could become active, if the preempt feature is utilized. This feature is useful in the following network situation. Consider that a router with a T1(1.544 Mbps) line is normally the preferred path from a LAN. If this router's serial interface should become inactive for any reason, HSRP can automatically change the active router by allowing the HSRP peer with the highest priority to assume the function of the active router. This can be used to ensure that the active router has connectivity to the rest of the network regardless of line problems. This feature also enable the usually preferred (T1 connected) router to regain its active HSRP role if connectivity is regained on its serial interface. This will occur automatically, because the priority of that router will increase with the change in the serial interface's line protocol. This will enable traffic to be rerouted though this preferred route transparently. In order to configure the tracking feature use this command "standby track [interface] [amount to decrement priority]":

It is possible that the use of the tracking feature can create a situation where an HSRP router gain an active forwarding role, due to the line protocol of a serial interface coming up, and not have a full routing table yet. Since this would be sub-optimal there is now a command designed to prevent this. This feature which was added in IOS 11.3 allows an administrator to indicate a configurable delay before the transition in active routers will take place. This allows time for the HSRP peer which has recently gained higher priority to populate its routing table before it takes on the active role. This is configured with the "preempt delay [interval]" command.

Figure 3



Syslog support:

Support for syslog messaging for HSRP information was added into IOS 11.3. This allows for more efficient logging and tracking of the current active and standby routers on syslog servers and it also enable monitoring of changes in HSRP state on multiple routers in one location. This feature allows for tracking of resources used in conjunction with M-HSRP also.

```
%STANDBY-6-STATECHANGE: Standby: 0: FastEthernet0.1 state Speak -> Standby
%STANDBY-6-STATECHANGE: Standby: 0: FastEthernet0.1 state Standby -> Active
```

The Ability to Use the BIA address

Beginning in 11.1.8, it is also possible for HSRP routers to use their burned-in MAC addresses (BIA) in conjunction with HSRP instead of only allowing the use of the HSRP MAC address, determined as a concatenation of the HSRP group number and the well known MAC prefix. Due to the fact that some Ethernet controllers on Cisco hardware are only able to use a single unicast MAC address filter, it was not previously possible for DECnet or XNS to be used in conjunction with HSRP on these platforms. This change enables the use of DECnet, XNS and HSRP on the same router by allowing the DECnet MAC address to be used as the HSRP MAC address. Using the BIA is also useful in networking situations where a device's BIA has been configured in other devices on the LAN, and in some configurations using source route bridging where a RIF might contain a particular MAC address.

However when using this feature only one standby group may be configured on that interface. When using "standby use-bia", the behavior of HSRP changes somewhat, such that the interface MAC address does not change when the router becomes active. Therefore, when the standby router takes over the MAC address associated with the HSRP IP address changes to the newly active router's MAC address. To enable non-disruptive fail-over when another router becomes active it sends out a gratuitous ARP to allow connected host devices to update their ARP cache entries.

FDDI Specific Command:

Due to an HSRP enhancement in IOS 11.3, all HSRP routers in a FDDI environment use their own unique burned-in MAC address to exchange messages and run the HSRP protocol. The active router will also claim the virtual MAC address. To ensure that learning bridges and switches cache the correct port entry for the virtual MAC address the active router also sends periodic messages using the HSRP MAC address. The rate at which these messages are sent can be configured by an administrator using the command "standby mac-refresh [interval]". Specifying an interval of 0 disables these MAC refresh messages.

Caveats

Ethernet Controller Type:

M-HSRP has the restriction that it can only be used in conjunction with certain Ethernet controllers used on Cisco IOS routers. This is the case because some Ethernet controllers can only support a single unicast MAC address. This restriction exists only on the lower end Cisco routers today. A matrix follows which details this restriction.

Interaction with Legacy Protocols:

There are further restrictions which exist with the use of the DECnet and XNS protocols, since they require the changing of the MAC addresses of all interfaces on the router in order to function. This can restrict the use of HSRP in this multi-protocol environment. If a router is configured for DECnet or XNS, all interfaces must share and use a specified MAC address. In order to use HSRP in conjunction with DECnet or XNS the command "standby use-bia" needs to be applied to an interface, as discussed earlier in this bulletin. This will dis-allow the use of M-HSRP. This will also have an impact on the DECnet or XNS functionality on that LAN.

The restrictions and limitations of the various platforms are shown in the following matrix:

Platform	100x	16xx	25xx	36xx	38xx	4xxx	52/53xx	58xx	RSM	7xxx	12000
HSRP	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
M-HSRP (Ethernet)	no	no	no	yes	yes	no	no	yes	yes	yes	no**
M-HSRP (other supported media***)	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes	yes
DECnet/XNS w HSRP	yes*	yes*	yes*	yes	yes	yes*	yes*	yes	yes	yes	no**
DECnet/XNS w M-HSRP	no	no	no	yes	yes	no	no	yes	yes	yes	no**

*with the "use bia" command.

** interface type or protocols not supported on this platform

***Fast Ethernet, token ring, FDDI, Gigabit Ethernet, ATM, including vLAN encapsulations: LANE, ISL, 802.10 SDE, 802.1Q, and FastEtherChannel

General Information

Addresses

In most cases when routers are configured to be part of a Hot Standby group, they will recognize their own native MAC address plus the hot standby group MAC address, unless they are only able to utilize one MAC address. Routers whose Ethernet controllers only recognize a single MAC address will use the HSRP MAC address when they are the active router and their burned in address (BIA) when they are in standby or not speaking, unless they are configured with the "use bia" command as explained above.

The following are the MAC addresses used by HSRP on all media which supports HSRP, except token ring:

0000.0c07.ac** (* - HSRP grp #)

Token Ring interfaces use functional addresses for the HSRP MAC addresses. Functional addresses are the only general multicast mechanism available. There are a limited number of token ring functional addresses and many of them are reserved for other functions. The following are the three addresses available for use with HSRP:

c000.0001.0000 (group 0)

c000.0002.0000 (group 1)

c000.0004.0000 (group 2)

You may only configure 3 HSRP addresses on token ring interfaces.

ICMP interaction:

The functionality of ICMP redirection is disabled automatically when using HSRP on an interface. The reason for this is that the possibility exists that hosts will cache entries resulting from these replies which may be sub-optimal in the case of a transition in the active router responsibilities in HSRP.

The use of Authentication with HSRP:

It is possible to use the authentication feature on HSRP packets sent to other routers in the HSRP group. The authentication built into HSRP consists of the addition of a shared clear-text key within the HSRP packets. The purpose of this password is to disallow mis-configured routers from participating in an HSRP group it was not intended to participate in. Although the feature allows for a certain level of security, in practice it is possible to cause more issues that this authentication can avert. Since the key is clear text on the wire, it would be relatively simple for an attacker bothering to forge an HSRP packet to sniff and replay this key, of course.

The behavior of an HSRP peer in the event that it receives a packet which contains a bad authentication value is to ignore it. This means that in the case of a mis-configured router which was meant to participate in an HSRP group it is possible that two or more routers will see no other valid hellos and decide to become active at the same time. In this event they would all be using the same IP and MAC address concurrently. This could create problems in the network. If the authentication feature is used on a network, it should be used cautiously. It is advisable to explicitly configure the standby address for which a router may assume responsibility in each device participating in an HSRP group.

Unsupported LAN interface types:

ATM BVI not supported in conjunction with HSRP at this time.

Future Direction of HSRP

Coming in future releases of IOS software are enhancements which will enable redundancy of other IOS features and innovations in conjunction with HSRP to further heighten the reliability and performance capabilities of your network. This will enable fault-tolerance for protocols and functionality which will increase the value of IOS throughout networks into the future.

Matrix of IOS release vs. HSRP functionality

The following is a matrix showing the support for different HSRP features in IOS:

Platform	10.0	10.2	10.3	11.0	11.1	11.2	11.3
HSRP	X	X	X	X	X	X	X
Standby Preempt	X	X	X	X	X	X	X
Ethernet 802.10SDE	--	--	--	--	X	X	X
Tracking	--	--	--	--	X	X	X
Use BIA	--	--	--	--	X ¹	X	X
Preempt Delay	--	--	--	--	--	X	X
Ethernet LANE	--	--	--	--	--	X	X
Inter Switch Link	--	--	--	--	--	--	X
Token Ring LANE	--	--	--	--	--	--	X
Syslog Support	--	--	--	--	--	--	X

1. in 11.1.8

References

draft-li-hsrp-01: Hot Standby Router Protocol (HSRP)

Cisco IOS Software Documentation: Internetworking Case Studies: Using HSRP for Fault-Tolerant IP Routing



Corporate Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters
Cisco Systems Europe s.a.r.l.
Parc Evolic, Batiment L1/L2
16 Avenue du Quebec
Villebon, BP 706
91961 Courtaboeuf Cedex
France
<http://www-europe.cisco.com>
Tel: 33 1 6918 61 00
Fax: 33 1 6928 83 26

Americas
Headquarters
Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
<http://www.cisco.com>
Tel: 408 526-7660
Fax: 408 527-0883

Asia Headquarters
Nihon Cisco Systems K.K.
Fuji Building, 9th Floor
3-2-3 Marunouchi
Chiyoda-ku, Tokyo 100
Japan
<http://www.cisco.com>
Tel: 81 3 5219 6250
Fax: 81 3 5219 6001

Cisco Systems has more than 200 offices in the following countries. Addresses, phone numbers, and fax numbers are listed on the

Cisco Connection Online Web site at <http://www.cisco.com>.

Argentina • Australia • Austria • Belgium • Brazil • Canada • Chile • China (PRC) • Colombia • Costa Rica • Czech Republic • Denmark
England • France • Germany • Greece • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia
Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Russia • Saudi Arabia • Scotland •
Singapore