

Hot Standby Router Protocol Features and Functionality

Document ID: 9234

Introduction

Prerequisites

- Requirements
- Components Used
- Conventions

HSRP Background and Operations

- Dynamic Router Discovery Mechanisms
- HSRP Operation

HSRP Addressing

Cisco IOS Release and HSRP Functionality Matrix

Cisco IOS Boot Images and HSRP Functionality

HSRP Features

- Preemption
- Interface Tracking
- Use Burned-In Address
- Multiple HSRP Groups
- Configurable MAC Address
- Syslog Support
- HSRP Debugging
- Enhanced HSRP Debugging
- Authentication
- IP Redundancy
- SNMP Management Information Base
- HSRP Support for Multiprotocol Label Switching Virtual Private Networks
- HSRP Support for ICMP Redirects

HSRP Interface and Media Support

- Ethernet
- Token Ring
- 802.1Q
- ISL
- FDDI
- MAC Refresh
- Bridge Group Virtual Interface
- Subinterfaces

Related Information

Introduction

This document describes the features and functionality of Hot Standby Router Protocol (HSRP).

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document is not restricted to specific software and hardware versions.

Conventions

For more information on document conventions, refer to Cisco Technical Tips Conventions.

HSRP Background and Operations

One way to achieve near-100 percent network uptime is to use HSRP, which provides network redundancy for IP networks, ensuring that user traffic immediately and transparently recovers from first hop failures in network edge devices or access circuits.

By sharing an IP address and a MAC (Layer 2) address, two or more routers can act as a single "virtual" router. The members of the virtual router group continually exchange status messages. This way, one router can assume the routing responsibility of another, should it go out of commission for either planned or unplanned reasons. Hosts continue to forward IP packets to a consistent IP and MAC address, and the changeover of devices doing the routing is transparent.

Dynamic Router Discovery Mechanisms

Below are descriptions of dynamic router discovery mechanisms that are available to hosts. Many of these mechanisms don't provide the network resiliency required by network administrators. This may be because the protocol wasn't initially intended to provide network resiliency or because it isn't feasible for every host on a network to be running the protocol. In addition to what is listed below, it is important to note that many hosts only allow you to configure a default-gateway.

Proxy Address Resolution Protocol

Some IP hosts use proxy Address Resolution Protocol (ARP) to select a router. When a host runs proxy ARP, it sends an ARP request for the IP address of the remote host it wants to contact. A router, Router A, on the network replies on behalf of the remote host and provides its own MAC address. With proxy ARP, the host behaves as if the remote host were connected to the same segment of the network. If Router A fails, the host continues to send packets destined for the remote host to the MAC address of Router A even though those packets have nowhere to go and are lost. You can either wait for ARP to acquire the MAC address of another router, Router B, on the local segment by sending another ARP request, or reboot the host to force it to send an ARP request. In either case, for a significant period of time, the host can't communicate with the remote host, even though the routing protocol has converged, and Router B is prepared to transfer packets that would otherwise go through Router A.

Dynamic Routing Protocol

Some IP hosts run (or snoop) a dynamic routing protocol such as the Routing Information Protocol (RIP) or Open Shortest Path First (OSPF) to discover routers. The drawback of using RIP is that it is slow to adapt to changes in the topology. Running a dynamic routing protocol on every host may not be feasible for a number of reasons, including administrative overhead, processing overhead, security issues, or lack of a protocol implementation for some platforms.

ICMP Router Discovery Protocol

Some newer IP hosts use ICMP Router Discovery Protocol (IRDP) (RFC 1256) 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. The default timer values of IRDP mean that it's not suitable for detection of failure of the first hop. The default advertisement rate is once every 7 to 10 minutes, and the default lifetime is 30 minutes.

Dynamic Host Configuration Protocol

Dynamic Host Configuration Protocol (DHCP) (RFC 1531) provides a mechanism for passing configuration information to hosts on a TCP/IP network. A host that runs a DHCP client requests configuration information from a DHCP server when it boots onto the network. This configuration information typically comprises an IP address and a default gateway. There is no mechanism for switching to an alternative router if the default gateway fails.

HSRP Operation

A large class of legacy host implementations that don't support dynamic discovery are capable of configuring a default router. Running a dynamic router discovery mechanism on every host may not be feasible for a number of reasons, including administrative overhead, processing overhead, security issues, or lack of a protocol implementation for some platforms. HSRP provides failover services to these hosts.

Using HSRP, a set of routers works in concert to present the illusion of a single virtual router to the hosts on the LAN. This set is known as an HSRP group or a standby group. A single router elected from the group is responsible for forwarding the packets that hosts send to the virtual router. This router is known as the Active router. Another router is elected as the Standby router. In the event that the Active router fails, the Standby assumes the packet-forwarding duties of the Active router. Although an arbitrary number of routers may run HSRP, only the Active router forwards the packets sent to the virtual router.

To minimize network traffic, only the Active and Standby routers send periodic HSRP messages once the protocol has completed the election process. If the Active router fails, the Standby router takes over as the Active router. If the Standby router fails or becomes the Active router, then another router is elected as the Standby router.

On a particular LAN, multiple hot standby groups may coexist and overlap. Each standby group emulates a single virtual router. The individual routers may participate in multiple groups. In this case, the router maintains separate state and timers for each group.

Each standby group has a single, well-known MAC address, as well as an IP address.

HSRP Addressing

In most cases when you configure routers to be part of an HSRP group, they listen for the HSRP MAC address for that group as well as their own burned-in MAC address. The exception is routers whose Ethernet controllers only recognize a single MAC address (for example, the Lance controller on the Cisco 2500 and Cisco 4500 routers). These routers use the HSRP MAC address when they are the Active router, and their burned-in address when they are not.

HSRP uses the following MAC address on all media except Token Ring:

0000.0c07.ac** (where ** is the HSRP group number)

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

```

c000.0001.0000 (group 0)
c000.0002.0000 (group 1)
c000.0004.0000 (group 2)

```

Note: When HSRP runs in a multiple–ring source–route bridging (SRB) environment and the HSRP routers reside on different rings, using the functional addresses can cause Routing Information Field (RIF) confusion. For example, in an SRB environment, it is possible that an HSRP standby router resides on a different ring than the active router. When this standby router becomes active, stations on the same ring as the old active router need a new RIF in order to send packets to the new active router. However, since the standby (new active) router is using the same functional address as the previous active router the stations are not aware that they must send explorers for a new RIF. For this reason, the **use–bia** command was introduced.

Cisco IOS Release and HSRP Functionality Matrix

This document shows which HSRP features are supported in which Cisco IOS® Software releases. Click on a feature to see a detailed description. An interim release number indicates in which release a feature first appeared, or a release where the functionality of that feature changed.

Feature	10.0	10.2	10.3	11.0	11.1	11.2	11.3	12.0	12.0T	12.1	12.1T
Preemption	X	X	X	X	X	X	X	X	X	X	X
Multiple Groups (MHSRP)			X	X	X	X	X	X	X	X	X
Ethernet 802.10 SDE					X	X	X	X	X	X	X
Interface Tracking					X	X	X	X	X	X	X
Use BIA					8.0	X	X	X	X	X	X
Preempt Delay						X	X	6.1	X	X	X
Ethernet LANE						X	X	X	X	X	X
Token Ring LANE							X	X	X	X	X
ISL							X	X	X	X	X
Syslog Support							X	X	X	X	X
MAC Refresh Interval								1.0	X	X	X
SNMP MIB									3.0	X	X
MHSRP and Use BIA									3.4	X	X
IP Redundancy									3.4	X	X
BVI									6.2	X	X

802.1Q									8.1	X	X
Enhanced HSRP Debugging										0.2	X
HSRP ICMP Redirects											3
HSRP MPLS VPNs											3

Cisco IOS Boot Images and HSRP Functionality

HSRP functionality was included in Cisco IOS boot images until the integration of Cisco bug ID CSCec16720 (registered customers only). Cisco bug ID CSCec16720 removed HSRP from boot images with the exception of:

- c7200-boot-mz
- c7200-kboot-mz
- c10k-eboot-mz
- c4500-boot-mz
- c7200-boot-mz
- c7200-kboot-mz
- c7400-kboot-mz
- ubr7200-boot-mz
- c6400r-boot-mz
- rpm-boot-mz
- rpmsf-boot-mz
- rsp-boot-mz
- urm-wboot-mz
- c5350-boot-mz
- c5400-boot-mz
- c7301-boot-mz
- c5850-boot-mz
- c4gwy-cboot-mz
- ubr910-rboot-mz
- ubr910-rboot-mz
- ubr925-k8boot-mz
- c5850tb-boot-mz

HSRP Features

Preemption

The HSRP preemption feature enables the router with highest priority to immediately become the Active router. Priority is determined first by the priority value that you configure, and then by the IP address. In each case a higher value is of greater priority.

When a higher priority router preempts a lower priority router, it sends a coup message. When a lower priority active router receives a coup message or hello message from a higher priority active router, it changes to the speak state and sends a resign message.

Preempt Delay

The preempt delay feature allows preemption to be delayed for a configurable time period, allowing the router to populate its routing table before becoming the active router.

Prior to Cisco IOS Software release 12.0(9), the delay started when the router reloaded. In Cisco IOS release 12.0(9) the delay starts when preemption is first attempted.

To configure HSRP priority and preemption, use the **standby [group] [priority number] [preempt [delay [minimum] seconds] [sync seconds]]** command.

Refer to the HSRP Documentation for more information on configuring HSRP.

Interface Tracking

Interface tracking allows you 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 the specified interface's line protocol goes down, the HSRP priority of this router is reduced, allowing another HSRP router with higher priority can become active (if it has preemption enabled).

To configure HSRP interface tracking, use the **standby [group] track interface [priority]** command.

When multiple tracked interfaces are down, the priority is reduced by a cumulative amount. If you explicitly set the decrement value, then the value is decreased by that amount if that interface is down, and decrements are cumulative. If you do not set an explicit decrement value, then the value is decreased by 10 for each interface that goes down, and decrements are cumulative.

The following example uses the following configuration, with the default decrement value of 10.

Note: When an HSRP group number is not specified, the default group number is group 0.

```
interface ethernet0
  ip address 10.1.1.1 255.255.255.0
  standby ip 10.1.1.3
  standby priority 110
  standby track serial0
  standby track serial1
```

The HSRP behavior with this configuration is:

- 0 interfaces down = no decrease (priority is 110)
- 1 interface down = decrease by 10 (priority becomes 100)
- 2 interfaces down = decrease by 10 (priority becomes 90)

The above HSRP behavior is true even if the decrement values are configured explicitly as below.

```
interface ethernet0
  ip address 10.1.1.1 255.255.255.0
  standby ip 10.1.1.3
  standby priority 110
  standby track serial0 10
  standby track serial1 10
```

Prior to Cisco IOS release 12.1, if you start a router with a down interface, HSRP interface tracking regards the interface as up.

This defect has Cisco bug ID CSCdp32289 (registered customers only) .

Use Burned-In Address

The use burned-in address (BIA) feature allows HSRP groups to use an interface's burned-in MAC address instead of an HSRP MAC address. Use BIA was first implemented in Cisco IOS release 11.1(8). To configure HSRP to use the BIA, use the **standby use-bia [scope interface]** command.

The **use-bia** command was implemented to overcome the limitations of using a functional address for the HSRP MAC address on Token Ring interfaces.

Note: When HSRP runs in a multiple-ring source-routed bridging environment and the HSRP routers reside on different rings, using the functional addresses can cause Routing Information Field (RIF) confusion. For this reason, the **use-bia** command was introduced.

The **use-bia** feature also enables the use of DECnet, Xerox Network Systems (XNS), and HSRP on the same router by allowing the DECnet MAC address (the BIA) to be used as the HSRP MAC address. The **use-bia** command is also useful in networking situations where a device's BIA has been configured in other devices on the LAN.

However, the **use-bia** command has several disadvantages:

- When a router becomes active, the virtual IP address is moved to a different MAC address. The newly active router sends a gratuitous ARP response, but not all host implementations handle the gratuitous ARP correctly.
- Proxy ARP breaks when **use-bia** is configured. A standby router cannot cover for the lost proxy ARP database of a failed router.
- Prior to Cisco IOS release 12.0(3.4)T, only one HSRP group is allowed if **use-bia** is configured.

When you configure the **use-bia** command on a subinterface, it actually shows up on the main interface and is applied to all subinterfaces. In Cisco IOS release 12.0(6.2) and later, the **use-bia** command is extended with the optional scope interface keywords to allow it to be applied to a single subinterface.

This defect has Cisco bug ID CSCdm25468 (registered customers only) .

Multiple HSRP Groups

The multiple HSRP (MHSRP) groups feature was added in Cisco IOS release 10.3. This feature further enables redundancy and load-sharing within networks, and 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 in the listen state for another group.

As of Cisco IOS release 12.0(3.4)T, you can use the **use-bia** command with multiple HSRP groups enabled.

Configurable MAC Address

Normally you use HSRP to help end stations locate the first hop gateway for IP routing. The end stations are configured with a default gateway. However, HSRP can provide first hop redundancy for other protocols. Some protocols, such as Advanced Peer-to-Peer Networking (APPN), use the MAC address to identify the first hop for routing purposes.

In this case, it is often necessary to be able to specify the virtual MAC address using the **standby mac-address** command. The virtual IP address is unimportant for these protocols. The actual syntax of the

command is **standby [group] mac-address mac-address**.

Note: You cannot use this command on a Token Ring interface.

Syslog Support

Support for syslog messaging for HSRP information was added in Cisco IOS release 11.3. This feature allows for more efficient logging and tracking of the current active and standby routers on syslog servers.

HSRP Debugging

Prior to Cisco IOS release 12.1, the HSRP debugging command was relatively simple. To enable HSRP debugging, you would simply use the **debug standby** command, which enabled output of HSRP state and packet information for all standby groups on all interfaces.

A debug condition was added in Cisco IOS release 12.0(2.1) that allows the output from the **standby debug** command to be filtered based upon interface and group number. The command utilizes the **debug condition** paradigm introduced in Cisco IOS release 12.0, as follows: **debug condition standby interface group**. The interface you specify must be a valid interface capable of supporting HSRP. The group can be any group (0 – 255).

You can set debug conditions for groups that do not exist, which allows you to capture debug information during the initialization of a new group.

You must enable **standby debug** order for any debug output to be produced. If you don't configure any **standby debug** conditions, then debug output is produced for all groups on all interfaces. If you configure at least one **standby debug** condition, then **standby debug** output is filtered according to all **standby debug** conditions.

Enhanced HSRP Debugging

Prior to Cisco IOS release 12.1(0.2), HSRP debugging was of limited use because information was lost in the noise from periodic hello messages. Thus the enhanced debugging feature was added in Cisco IOS 12.1(0.2).

The following table explains the command options for enhanced debugging.

Command	Description
debug standby	Displays all HSRP errors, events, and packets.
debug standby terse	Displays all HSRP errors, events, and packets, except hello and advertisement packets.
debug standby errors	Displays HSRP errors.
debug standby events [[all terse] [icmp protocol redundancy track]] [detail]	Displays HSRP events.
debug standby packets [[all terse] [advertise coup hello resign]] [detail]	Displays HSRP packets.

You can filter the **debug** output using interface and HSRP group conditional debugging. To enable interface conditional debugging, use the **debug condition interface** *interface* command. To enable HSRP conditional debugging, use the **debug condition standby** *interface group* command.

An interface debug condition applies only when you have set no **standby debug** conditions. HSRP debugging is further enhanced in Cisco IOS release 12.1(1.3), based on the improvements that were made to the HSRP state table.

This defect has Cisco bug ID CSCdp57811 (registered customers only).

These enhancements display the HSRP state table events. In the output below the **a/**, **b/**, **c/**, and so on, refer to the events of the HSRP finite state machine, which are documented in RFC 2281 .

```
SB1: Ethernet0/2 Init: a/HSRP enabled
SB1: Ethernet0/2 Active: b/HSRP disabled (interface down)
SB1: Ethernet0/2 Listen: c/Active timer expired (unknown)
SB1: Ethernet0/2 Active: d/Standby timer expired (20.0.0.3)
SB1: Ethernet0/2 Speak: f/Hello rcvd from higher pri Speak router
SB1: Ethernet0/2 Active: g/Hello rcvd from higher pri Active router
SB1: Ethernet0/2 Speak: h/Hello rcvd from lower pri Active router
SB1: Ethernet0/2 Standby: i/Resign rcvd
SB1: Ethernet0/2 Active: j/Coup rcvd from higher pri router
SB1: Ethernet0/2 Standby: k/Hello rcvd from higher pri Standby router
SB1: Ethernet0/2 Standby: l/Hello rcvd from lower pri Standby router
SB1: Ethernet0/2 Active: m/Standby mac address changed
SB1: Ethernet0/2 Active: n/Standby IP address configured
```

Authentication

The HSRP authentication feature consists of a shared clear-text key contained within the HSRP packets. This feature prevents the lower priority router from learning the standby IP address and standby timer values from the higher priority router.

To configure the HSRP authentication string, use the **standby authentication** *string* command.

IP Redundancy

HSRP provides stateless redundancy for IP routing. HSRP by itself is limited to maintaining its own state. It assumes that each router builds and maintains its own routing tables independently of other routers. The IP redundancy feature provides a mechanism that allows HSRP to provide a service to client applications so that they can implement statefull failover.

IP redundancy does not provide a mechanism for peer applications to exchange state information. This is left to the applications themselves, and is essential if the applications are to provide statefull failover.

IP redundancy is currently (as of January 2000) implemented only for Mobile IP Home Agents. Following is a sample configuration:

```
configure terminal
router mobile
ip mobile home-agent standby hsrp-group1
!
interface e0/2
no shutdown
ip address 20.0.0.1 255.0.0.0
standby 1 ip 20.0.0.11
standby 1 name hsrp-group1
```

Note: As of Cisco IOS release 12.1(3)T, the keyword **redundancy** is accepted in addition to the keyword **standby**. The **standby** keyword will be phased out in a later Cisco IOS release. The correct command will then be **ip mobile home-agent redundancy hsrp-group1**.

Future uses of IP Redundancy may include:

- NAT – Need to provide redundant gateways.
- IPSEC – Need to synchronize state information in order to operate when HSRP is in use.
- DHCP Server – DHCP servers implemented in various routers.
- NBAR, CBAC – Need to mirror firewall states for asymmetric routing.
- GPRS – Needs a way to track TCP state.
- PIX

SNMP Management Information Base

SNMP Management Information Base (MIB) support was added to Cisco IOS release 12.0(3.0)T. There are two relevant MIBs for HSRP:

- ciscoMgmt 106: The MIB module for managing HSRP
- ciscoMgmt 107: The extension MIB module for managing HSRP

Prior to Cisco IOS release 12.0(6.1)T, a walk of the extended HSRP MIB when a Bridge Group Virtual Interface (BVI) is present causes the router to crash.

This defect has Cisco bug ID CSCdm61257 (registered customers only) .


HSRP Support for Multiprotocol Label Switching Virtual Private Networks

HSRP support for Multiprotocol Label Switching Virtual Private Networks (MPLS VPNs) was added in Cisco IOS release 12.1(3)T.

HSRP on an MPLS VPN interface is useful when you have an Ethernet connected between two Provider Edges (PEs) and you have either of the following:

- A Customer Edge (CE) with a default route to the HSRP virtual IP address.
- One or more hosts with the HSRP virtual IP address configured as the default gateway.

The network diagram below shows two PEs with HSRP running between their VPN routing/forwarding (VRF) interfaces. We configured the CE with the HSRP virtual IP address as its default route. And we configured HSRP to track the interfaces connecting the PEs to the rest of the provider network. For example, if interface E1 of PE1 fails, the HSRP priority will be reduced such that PE2 takes over forwarding packets to the virtual IP/MAC address.



Below are the configurations.

Router PE1	Router PE2
<pre> conf terminal ! ip cef ! ip vrf vrf1 rd 100:1 route-target export 100:1 route-target import 100:1 ! interface ethernet0 no shutdown ip vrf forwarding vrf1 ip address 10.2.0.1 255.255.0.0 standby 1 ip 10.2.0.20 standby 1 priority 105 standby 1 preempt delay minimum 10 standby 1 timers 3 10 standby 1 track ethernet1 10 standby 1 track ethernet2 10 </pre>	<pre> conf terminal ! ip cef ! ip vrf vrf1 rd 100:1 route-target export 100:1 route-target import 100:1 ! interface ethernet0 no shutdown </pre>

You can use the following commands to verify that the HSRP virtual IP address is in the correct VRF ARP and Cisco Express Forwarding tables:

```

ed1-pel1# show ip arp vrf vrf1
Protocol Address Age (min) Hardware Addr Type Interface
Internet 10.2.0.1 0 00d0.1bd3.bc22 ARPA Ethernet0/2
Internet 10.2.0.20 0 0000.0c07.ac01 ARPA Ethernet0/2

ed1-pel1# show ip cef vrf vrf1
Prefix Next Hop Interface
0.0.0.0/0 10.3.0.4 Ethernet0/3
0.0.0.0/32 receive

```

10.1.0.0/16	10.2.0.1	Ethernet0/2
10.2.0.0/16	attached	Ethernet0/2
10.2.0.1/32	receive	
10.2.0.20/32	receive	
224.0.0.0/24	receive	
255.255.255.255/32	receive	

HSRP Support for ICMP Redirects

HSRP is based on the concept that the HSRP peer routers protecting a subnet can provide access to all other subnets that comprise the network. Therefore, it is irrelevant which router becomes the active HSRP router, as all routers had routes to every subnet.

HSRP makes use of a special virtual IP address and virtual MAC address, which are logically attached to the HSRP active router. ICMP redirects are automatically disabled on an interface when using HSRP on that interface. IOS 12.1(3)T onwards, ICMP Redirects feature enables ICMP redirects on interfaces configured with HSRP. Refer to HSRP Support for ICMP Redirects for more details. This is done to prevent hosts from being redirected away from the HSRP virtual IP address. It is possible that the two (or more) routers on a subnet don't have identical connectivity to the rest of the network. That is, for a particular destination IP address, one or the other of the routers may have a much better path to that address, or may even be the only router attached to that address.

The ICMP protocol allows a router to redirect an endstation to send packets for a particular destination to another router on the same subnet, if the first router knows that the other router has a better path to that particular destination. As was the case for default gateways, if the router to which an endstation has been redirected for a particular destination fails, then the endstation's packets to that destination were not delivered. In standard HSRP, this is exactly what happens. For this reason, we recommend disabling ICMP redirects if HSRP is turned on.

Extending the relationship between ICMP redirects and HSRP provides a solution to this problem, allowing you to take advantage of the benefits of both HSRP and ICMP redirects. Two (or more) HSRP groups are run on each subnet, with at least as many HSRP groups configured as there are routers participating. The priorities are configured so that each of the routers is master of at least one HSRP group. When one router determines to redirect an endstation to a different router for a specific destination, then instead of redirecting the endstation to that other router's IP address, it finds an HSRP group that is being mastered by that router, and redirects the endstation to the corresponding virtual IP address. If that target router then fails, HSRP ensures that another router takes over its job and, perhaps, redirects the endstation to yet another, again virtual, router.

HSRP Interface and Media Support

This section explains which interfaces and media HSRP supports, and possible caveats when running HSRP over these media.

Since Cisco IOS Software release 10.0, HSRP functionality has been available on Ethernet, Token Ring and Fiber Distributed Data Interface (FDDI). Fast Ethernet and ATM interfaces are also supported by HSRP.

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. HSRP VLAN support was added in Cisco IOS release 11.1 for IEEE 802.10 Secure Data Exchange (SDE), and in Cisco IOS release 11.3 for Cisco Inter-Switch Link (ISL).

Ethernet

Several Ethernet (Lance and QUICC) controllers in low-end products can only have a single unicast MAC address in their address filter. On these platforms only a single HSRP group is permitted, and the interface address is changed to the HSRP virtual MAC address when the group becomes Active. If you are using HSRP on routers with multiple interfaces of this type, you should configure each interface with a different HSRP group number.

Note: The Cisco 7200 router also uses the Lance Ethernet controller, but it supports MHSRP in software.

Cisco recommends that you have no more than twenty-four HSRP Ethernet Interface Processors (EIPs) due to the time it takes to update the address filters for HSRP. Having more than twenty-four HSRP EIPs can cause instability and excessive CPU load.

This defect has Cisco bug ID CSCdj29595 (registered customers only) .

If you have more than twenty-four EIPs, try replacing the EIPs with Versatile Interface Processors (VIPs) and Ethernet port adapters. VIPs have been approved for up to eighty HSRP groups. You can also reduce the number of HSRP groups, and increase the HSRP hello and hold time.

Token Ring

One limitation of running HSRP on a Token Ring interface is that you cannot reprogram the address filter on the Token Ring chipset the same way you can on Ethernet, FDDI or ATM emulation. Token Ring uses functional addresses, of which there are only a small number available that do not conflict with other uses of the functional address space.

When running HSRP in a source-route bridging (SRB) environment, the use of functional addresses can cause RIF confusion. See the HSRP Addressing section for more information. Also, try configuring the **use-bia** command.

802.1Q

Cisco recommends using Cisco IOS release 12.0(8.1)T or later for HSRP over 802.1Q.

ISL

HSRP over ISL is available in Cisco IOS releases 11.2(6)F, 11.3, 12.X. It is recommended to use release 12.0(7) or later in order to avoid the issue described in Cisco bug ID CSCdm68811 (registered customers only) .

FDDI

A FDDI port adapter strips frames from the ring if it sees one of its own MAC addresses in the MAC source. If a network event causes both routers to go active, then both routers send HSRP hello packets with the same virtual MAC address. Each router mistakenly strips the other router's hello packet from the network, and both stay active.

This defect has Cisco bug ID CSCdj30049 (registered customers only) .

The solution to this problem in Cisco IOS release 11.2(11.1) is for HSRP routers in an FDDI environment to use their own unique burned-in MAC address to exchange messages and run the HSRP protocol. To ensure that learning bridges and switches cache the correct port entry for the virtual MAC address, the active router also sends periodic refresh messages using the HSRP MAC address.

Note: The Cisco 4500 router's hardware content-addressable memory (CAM) on an FDDI interface may not be populated correctly after a reload if you have configured multiple RIP networks and HSRP groups. The only workaround at this time is to clear the interfaces to restore the CAM. This defect has Cisco bug ID CSCdm93122 (registered customers only) .

MAC Refresh

HSRP routers in an FDDI environment use their own unique burned-in MAC address to exchange messages and run the HSRP protocol. To ensure that learning bridges and switches cache the correct port entry for the virtual MAC address the active router also sends periodic refresh messages using the HSRP MAC address. This defect has Cisco bug ID CSCdj30049 (registered customers only) .

If you do not have a switch or learning bridge on your network, you can disable the sending of refresh packets as shown below:

```
interface fddi 1/0/0
 ip address 10.1.1.1 255.255.255.0
 standby ip 10.1.1.250
 standby mac-refresh 0
```

Bridge Group Virtual Interface

HSRP support for Bridge Group Virtual Interfaces (BVI) was added in Cisco IOS release 12.0(6.2)T.

Subinterfaces

HSRP groups on subinterfaces must have a group number unique among all other groups on all subinterfaces on the same main interface. This is because subinterfaces do not receive a unique SNMP interface index. If you had two groups with the number N on different subinterfaces, then in the MIB, group N on sub-interface 1 and group N on sub-interface 2 would appear to be the same group.

Related Information

- [HSRP Support Page](#)
- [Technical Support & Documentation – Cisco Systems](#)

[Contacts & Feedback](#) | [Help](#) | [Site Map](#)

© 2009 – 2010 Cisco Systems, Inc. All rights reserved. [Terms & Conditions](#) | [Privacy Statement](#) | [Cookie Policy](#) | [Trademarks of Cisco Systems, Inc.](#)

Updated: May 25, 2006

Document ID: 9234
