Understand the Hot Standby Router Protocol Features and Functionality

Contents

Introduction

This document describes how the Hot Standby Router Protocol (HSRP) functions and reviews its features.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

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

The information in this document was created from the devices in a specific lab environment. All of the devices used in this document started with a cleared (default) configuration. If your network is live, ensure that you understand the potential impact of any command.

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, and ensures that user traffic immediately and transparently recovers from first hop failures in network edge devices or access circuits.

When two or more routers share an IP address and a MAC (Layer 2) address, they 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 if one is 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 that do the routing is transparent.

Dynamic Router Discovery Mechanisms

This provides descriptions of dynamic router discovery mechanisms that are available to hosts. Many of these mechanisms do not provide the network resiliency required by network administrators. This can be caused when the protocol was not initially intended to provide network resiliency or because it is not feasible for every host on a network to run the protocol. In addition to what is listed, it is important to notice 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 that sends another ARP request or reboots the host to force it to send an ARP request. In either case, for a significant period of time, the host cannot 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 when you use RIP is that it is slow to adapt to changes in the topology. To run a dynamic routing protocol on every host, this is not practical for a number of reasons, along with administrative overhead, processing overhead, security issues, or lack of a protocol implementation for some platforms considerations.

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 is 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 to pass the 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 does not support dynamic discovery can configure a default router. To run a dynamic router discovery mechanism on every host, is not practical for a number of reasons, along with administrative overhead, processing overhead, security issues, or lack of a protocol implementation for some platforms considerations. HSRP provides failover services to these hosts.

When you use 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 distribution of 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 can 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 can coexist and overlap. Each standby group emulates a single virtual router. The individual routers can 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 this MAC address on all media except Token Ring:

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

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	**12.0**	**12.0T**	**12.1**	**12.1T**
Preemption	X	X	X	X
Multiple Groups (MHSRP)	X	X	X	X
Ethernet 802.10 SDE	X	X	X	X
Interface Tracking	X	X	X	X
Use BIA	X	X	X	X
Preempt Delay	6.1	X	X	X
Ethernet LANE	X	X	X	X
Token Ring LANE	X	X	X	X
ISL	X	X	X	X
Syslog Support	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 HSRP Functionality

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, and allows the router to populate its routing table before it becomes the active router.

Before 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> <prioritynumber> <preempt [delay [minimum]seconds] [syncseconds]>** command. Refer to the HSRP Documentation for more information.

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 line protocol of the interface goes down, the HSRP priority of this router is reduced, and allows 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.

Note: The Interface Track command availability can depend on the software version used, but the **standby <group> track <object>** command can be used instead.

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.

This example uses this 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 becomes100)
- 2 interfaces down = decrease by 10 (priority becomes 90)

The previous mentioned HSRP behavior is true even if the decrement values are configured explicitly as:

```
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
```

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

Use Burned-In Address

The use burned-in address (BIA) feature allows HSRP groups to use a burned-in MAC address of the interface 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 when a functional address for the HSRP MAC address on Token Ring interfaces is used.

Note: When HSRP runs in a multiple-ring source-routed bridging environment and the HSRP routers reside on different rings and use the functional addresses, it 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 the use of the DECnet MAC address (the BIA) to be used as the HSRP MAC address. The **use-bia** command is also useful for networks where BIA of the device 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 sub-interface, 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 sub-interface.

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 forwards 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. Refer to Load Sharing with HSRP to configure HSRP and take advantage of multiple paths.

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 that uses 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

Before 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 that can support 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 do not 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 by all **standby debug** conditions.

Enhanced HSRP Debugging

Before 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 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 with the 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 software release 12.1(1.3), based on the improvements that were made to the HSRP state table.

These enhancements display the HSRP state table events. In the output, 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 (10.0.0.3)
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 can maintain only 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 stateful 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 stateful failover.

IP redundancy is commonly implemented only for Mobile IP Home Agents. This is a sample configuration:

```
configure terminal
 router mobile
 ip mobile home-agent standby hsrp-group1
!
interface e0/2
 no shutdown
 ip address 10.0.0.1 255.0.0.0
 standby 1 ip 10.0.0.11
 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 is phased out in a later Cisco IOS release. The correct command is **ip mobile home-agent redundancy** **hsrp-group1** .

Future uses of IP Redundancy 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.

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 used to manage HSRP
- ciscoMgmt 107: The extension MIB module used to manage 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.

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 these:

- 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 shows two PEs with HSRP that run between their VPN routing/forwarding (VRF) interfaces. The CE with the HSRP virtual IP address are configured as its default route. And HSRP is configured to track the interfaces that connect the PEs to the rest of the provider network. For example, if interface E1 of PE1 fails, the HSRP priority is reduced such that PE2 takes over forwarding packets to the virtual IP/MAC address.

These are the configurations:

Router PE1	Router PE2
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	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 10.2.0.2 255.255.0.0 standby 1 ip 10.2.0.20 standby 1 priority 100 standby 1 preempt delay minimum 10

standby 1 timers 3 10 standby 1 track ethernet1 10 standby 1 track ethernet2 10	standby 1 timers 3 10 standby 1 track ethernet1 10 standby 1 track ethernet2 10

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

```
<#root>

ed1-pe1#

show ip arp vrf vrf1

Protocol  Address          Age (min)  Hardware Addr   Type   Interface
Internet  10.2.0.1                -   00d0.bbd3.bc22  ARPA   Ethernet0/2
Internet  10.2.0.20               -   0000.0c07.ac01  ARPA   Ethernet0/2

ed1-pe1#

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 that protect 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 HSRP is used on that interface. Cisco 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 redirection away from the HSRP virtual IP address. It is possible that the two (or more) routers on a subnet do not have identical connectivity to the rest of the network. That is, for a particular destination IP address, one or the other of the routers can have a much better path to that address or can 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 packets to that destination were not delivered. In standard HSRP, this is exactly what happens. For this reason, it is recommended that you disable ICMP redirects if HSRP is turned on.

When you extend the relationship between ICMP redirects and HSRP provides a solution to this problem

and this allows 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 that participate. The priorities are configured so that each router is the primary router for at least one HSRP group. When one router determines to redirect an endstation to a different router for a specific destination, then instead of the redirect to the endstation to that other router IP address, it finds an HSRP group that has that router as its primary 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 virtual router.

Bridge Group Virtual Interface

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

Sub-interfaces

HSRP groups on sub-interfaces must have a group number unique among all other groups on all sub-interfaces on the same main interface. This is because sub-interfaces do not receive a unique SNMP interface index. If you had two groups with the number N on different sub-interfaces, 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**
- **HSRP - FAQ**