Review Hot Standby Router Protocol (HSRP) FAQ

Contents

Introduction

This document describes the most frequently asked questions about and related to the Hot Standby Router Protocol (HSRP).

Q. Can the standby router take over if the active router LAN interface state is "interface up line protocol down"?

A. Yes, the standby router takes over once the hold time expires. By default, this is equivalent to three hello packets from the active router having been missed. The actual convergence time depends on the HSRP timers configured for the group and possibly on routing protocol convergence. The HSRP hello time timer defaults to three and the hold time timer defaults to ten.

Q. Can I configure more than one standby group with the same group number?

A. Yes. However, Cisco does not recommend it on lower-end platforms such as the 4x00 series and earlier. If the same group number is assigned to multiple standby groups, it creates a non-unique MAC address. This is seen as the MAC address of the router, and it is filtered out if more than one router in a LAN becomes active. This behavior can change in future releases of Cisco IOS®.

> **Note**: The 4x00 series and earlier do not have the hardware required to support more than one MAC address at a time on Ethernet interfaces. However, the Cisco 3600 and newer platforms do support multiple MAC addresses on all Ethernet interfaces.

Q. When an active router tracks serial 0 and the serial line goes down, how does the standby router know to become active?

A. When the state of a tracked interface changes to down, the active router decrements its priority. The standby router reads this value from the hello packet priority field and becomes active if this value is lower than its own priority and the standby preempt is configured. You can configure by how much the router must decrement the priority. By default, it decrements its priority by ten.

Q. If there is no priority configured for a standby group, what determines which router is active?

A. The priority field is used to elect the active router and the standby router for the specific group. In the case of an equal priority, the router with the highest IP address for the respective group is elected as active. Furthermore, if there are more than two routers in the group, the second highest IP address determines the standby router and the other router/routers are in the listen state.

> **Note**: If no priority is configured, it uses the default of 100.

Q. What are the limiting factors that determine how many standby

groups can be assigned to a router?

A. Ethernet: 256 per router. **FDDI**: 256 per router. **Token Ring**: 3 per router (uses reserved functional address).

Q. Which HSRP router requires that I configure preempt?

A. An HSRP-enabled router with preempt configured attempts to assume control as the active router when its Hot Standby priority is higher than the current active router. The **standby preempt** command is needed in situations when you want an occurring state change of a tracked interface to cause a standby router to take over from the active router. For example, an active router tracks another interface and decrements its priority when that interface goes down. The standby router priority is now higher and it sees the state change in the hello packet priority field. If preempt is not configured, it cannot take over and failover does not occur.

Q. Based on the documentation, it looks like I can use HSRP to achieve load-balancing across two serial links. Is this true?

A. Yes, refer to [Load Sharing with HSRP](#) for more information.

Q. Does HSRP support DDR, and if so, how can it know to dial?

A. No, HSRP does not support Dial-on-Demand Routing (DDR) directly. However, you can configure it to track a serial interface and swap from the active to the standby router in case of a WAN link failure. The command used to track the state of an interface is standby <group#> track <interface> .

Q. I use HSRP and all hosts use the active router to forward traffic to the rest of my network. I have noticed that the return traffic comes back through the standby router. Can this cause problems with HSRP or my applications?

A. No, normally this is transparent to all hosts and/or servers on the LAN and can be desirable if a router experiences high traffic. In order to change this, configure a more desirable cost for the link you want the distant router/routers to use.

Q. Can two Cisco routers of different model on the same LAN segment use HSRP, or do I have to replace one of the routers so the platforms are identical?

A. You can mix the platforms with HSRP, but you are not able to support multiple HSRP (MHSRP) due to the hardware limitations of the lower-end platform.

Q. If I use a switch, what do I see on the CAM tables for the HSRP?

A. The content-addressable memory (CAM) tables provide a map for the HSRP MAC address to the port on which the active router is located. In this way, you can determine what the switch perceives the HSRP status to be.

Q. What is the standby use-bia command and how does it work?

A. By default, HSRP uses the preassigned HSRP virtual MAC address on Ethernet and FDDI, or the functional address on Token Ring. In order to configure HSRP to use the burnt-in address of the interface as its virtual MAC address, instead of the default, use the standby use-bia command.

For example, on Token Ring, if Source Route Bridging is in use, a Routing Information Field (RIF) is stored with the virtual MAC address in the RIF cache of the host. The RIF indicates the path and final ring used to reach the MAC address. As routers transition to the active state, they send gratuitous Address Resolution Protocols (ARPs) in order to update the ARP table of the host. However, this does not affect the RIF cache of the hosts that are on the bridged ring. This situation can lead to packets being bridged to the ring for the previous active router. In order to avoid this situation, use the standby use-bia command. The router now uses its burnt-in MAC address as the virtual MAC address.

The use of standby use-bia command has these 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 the failed router.

Q. Can I run NAT and HSRP together?

A. You can configure network address translation (NAT) and HSRP on the same router. However, a router that runs NAT holds state information for traffic that is translated through it. If this is the active HSRP router and the HSRP standby takes over, the state information is lost.

> **Note**: Stateful NAT (SNAT) can make use of HSRP to fail over. Refer to NAT Stateful Failover of Network Address Translation for more information. Static NAT Mapping Support with HSRP for High Availability is another feature which makes NAT and HSRP interact. If static NAT is configured with the same IP on each router, the routers advertise each other with the MAC addresses, and the routers display the %IP-4-DUPADDR: Duplicate address [ip address] on [interface], sourced by [mac-address] error message. Refer to NATâ€"Static Mapping Support with HSRP for High Availability for more information.

> **Note**: SNAT End-of-Life was announced and was stopped January, 2012. For more information, refer to End-of-Sale and End-of-Life Announcement for the Cisco IOS Stateful Failover of Network Address Translation (SNAT).

Q. What are the IP source address and destination address of HSRP hello packets?

A. The destination address of HSRP hello packets is all routers multicast address (224.0.0.2). The source address is the primary IP address of the router assigned to the interface.

Q. Are HSRP messages TCP or UDP?

A. UDP, since HSRP runs on UDP port 1985.

Q. HSRP does not work when an Access Control List (ACL) is applied. How can I permit HSRP through an ACL?

A. HSRP hello packets are sent to multicast address 224.0.0.2 with UDP port 1985. Whenever an ACL is applied to an HSRP interface, ensure that packets destined to 224.0.0.2 on UDP port 1985 are permitted.

Q. How does TACACS/RADIUS accounting work with HA routers with HSRP?

A. If routers are configured in HA mode (that run HSRP in-between them), then the active and standby routers act as one logical unit and share the same IP and MAC address. Only the active router generates the accounting record with a particular virtual IP address and updates the TACACS/RADIUS server. If the standby generates the accounting record with the same address, there is duplicate data in the backend RADIUS/TACACS server. Therefore, in order to avoid duplication of data, the standby router does not generate accounting records.

Q. Are HSRP and VLAN translation supported together in a Cisco Catalyst 6500 series switch?

A. VLAN translation and HSRP can be configured together in a Cisco Catalyst 6500 series switch, subject to the restrictions put in place by VLAN translation. Refer to [VLAN Translation Guidelines and Restrictions](#) for more information.

Q. Is it possible to use HSRP to track the tunnel interface?

A. It is not possible to use the HSRP configuration to track the GRE tunnel interface. However, the tunnel interface never goes down and the track never triggers failover.

Q. How do I perform a forced failover of an HSRP active router without a shutdown on an interface?

A. The only way to make a failover without an interface shut down is to manually change the priority in the HSRP configuration.

Q. Is it possible to run HSRP on an interface configured for 802.1q trunking?

A. Yes it is possible to run HSRP on the interfaces configured for 802.1q. Make sure to verify that both sides of the trunk are configured to use the same native VLAN and verify that VLANs are not pruned and in the STP state for router-connected ports.

Q. Is it possible to run HSRP between two routers on two different

interfaces?

A. Yes, it is possible to run HSRP on two interfaces on two different routers. In order to have HSRP on two interfaces on two different routers, two HSRP groups are needed.

Q. Is it possible to run HSRP and OSPF together on a backbone router?

A. Yes, however HSRP and OSPF are two different protocols that achieve different things. The OSPF that runs on the router advertises the two physical interfaces and not the virtual IP address. When this router becomes active, it broadcasts a gratuitous ARP packet with the HSRP virtual MAC address to the affected LAN segment. If the segment uses an Ethernet switch, this allows the switch to change the location of the virtual MAC address so that packets go to the new router instead of the one that is no longer active. End devices do not actually need this gratuitous ARP if the routers use the default HSRP MAC address.

Q. Which IP address must be seen when a reply is received for traceroute?

A. When a reply for traceroute is received from a hop that runs HSRP, the reply must contain the active physical IP adddress and not the virtual ip address. If there is an asymmetric routing in the network due to which standby router IP address is seen in the reply for the traceroute.

Q. What is the difference between GLBP and HSRP?

A. GLBP provides load balancing over multiple routers (gateways) using a single virtual IP address and multiple virtual MAC addresses. Members of the GLBP group select one of them to become the active virtual gateway for the group.

With HSRP in a single router (gateway), one interface is used as the active interface and the other interface is in standby. The active interface is used for all the traffic and the standby interface just waits for the active interface to fail without any traffic.

Q. Is it possible to run HSRP for both primary and secondary subnets?

A. Yes. The use of HSRP for secondary addresses is supported. This feature along with the multiple HSRP feature is beneficial in real networks. Refer to the Multiple HSRP Groups & Secondary Addresses section of [Understand the Hot Standby Router Protocol Features and Functionality](#) for the configuration example.

Q. What is the use of delay in standby preempt delay minimum 60 command?

A. If router A is the HSRP active router and then loses a link, which causes it to become standby router, and then the link comes back, the `delay` command causes router A to wait before it becomes active again. In this case, it waits for 60 seconds for the router to become active.

Q. Is it possible to run HSRP on a sub-interfaces?

A. Yes. You can run HSRP on Sub-Interfaces.

Q. Is it possible to track specific routes with specific outgoing interface in HSRP?

A. Tracking a particular route is an option. When a particular route is unavailable, track goes down. Based on that track, you can configure HSRP to switchover. Use this configuration:

```
track 10 ip sla 123 reachability
 delay down 10 up 10
!
ip sla 123
 icmp-echo <A.B.C.D> timeout 20000
!
ip sla schedule 123 life forever start-time now

### To call this track in hsrp ###

interface <interface name>
 standby 1 track 10
```

Q. I get the % Warning: Interface MAC address filter only supports 28 additional addresses % and 28 HSRP groups are already configured. The HSRP MAC address cannot be % added to the MAC address filter if the group becomes active. error message when the HSRP group is configured under the portchannel interface. Why?

A. This error message appears due to the hardware limitation of platform. There are 28 HSRP-groups that can be supported by a Port-channel interface.

Q. How do you track default route reachability in GLBP?

A. Use this configuration:

```
track 10 ip route 0.0.0.0 0.0.0.0 reachability
!
interface fa0/1
glbp 50 ip <ip address>
glbp 50 priority 210
glbp 50 preempt
glbp 50 weighting track 10
```

Q. What are the differences between HSRP version 2 and HSRP

version 1?

A. The differences between HSRP version 2 and HSRP version 1 are as follows:

- In HSRP version 1, millisecond timer values are not advertised or learned. HSRP version 2 advertises and learns millisecond timer values. This change ensures stability of the HSRP groups in all cases.
- The group numbers in version 1 are restricted to the range from 0 to 255. HSRP version 2 expands the group number range from 0 to 4095. For example, new MAC address range can be used, 0000.0C9F.Fyyy, where yyy = 000-FFF (0-4095).
- HSRP version 2 uses the new IP multicast address 224.0.0.102 to send hello packets instead of the multicast address of 224.0.0.2, which is used by version 1.
- HSRP version 2 packet format includes a 6-byte identifier field that is used to uniquely identify the sender of the message. Typically, this field is populated with the interface MAC address. This improves troubleshooting network loops and configuration errors.
- HSRP version 2 allows for future support of IPv6.
- HSRP version 2 has a different packet format than HSRP version 1. The packet format uses a type-length-value (TLV) format. HSRP version 2 packets received by an HSRP version 1 router can have the type of field mapped to the version field by HSRP version 1, and subsequently ignored.
- A new command can allow changing of the HSRP version on a per-interface level standby version [1 | 2]. Note that HSRP version 2 cannot interoperate with HSRP version 1. However, the different versions can be run on different physical interfaces of the same router.

Q. Can I configure HSRP on Catalyst 9300 Series Switches?

A. Yes, HSRP can be configured on Catalyst 9300 Series Switches. In order to view sample configuration commands, refer to Configuring HSRP.

> **Note**: Use the Cisco Feature Navigator tool in order to verify the HSRP support on the Cisco IOS image.

Related Information

- **HSRP Support**
- **Hot Standby Router Protocol Features and Functionality**
- **HSRP Support Page**
- **Cisco Technical Support & Downloads**