Troubleshoot HSRP Common Issues

Contents

- **Introduction**
- **Prerequisites**
  - Requirements
  - Components Used
- **Understand HSRP**
- **Background Information**
  - Basic Operation
  - HSRP Terms
  - HSRP Addressing
    - HSRP Router Communication
    - HSRP Standby IP Address Communication on All Media Except Token Ring
  - ICMP Redirects
  - HSRP Functionality Matrix
  - HSRP Features
  - Packet Format
  - HSRP States
  - HSRP Timers
  - HSRP Events
  - HSRP Actions
  - HSRP State Table
  - Packet Flow
    - Router A Configuration (Active Router)
    - Router B Configuration (Standby Router)
- **Troubleshoot HSRP Case Studies**
  - Case Study #1: HSRP Standby IP Address Is Reported as a Duplicate IP Address
  - Case Study #2: HSRP State Continuously Changes (Active, Standby, Speak) or %HSRP-6-STATECHANGE
  - Case Study #3: HSRP Does Not Recognize Peer
  - Case Study #4: HSRP State Changes and Switch Reports SYS-4-P2_WARN: 1/Host Is Flapping Between Port and Port in Syslog
  - Case Study #5: Asymmetric Routing and HSRP (Excessive Flooding of Unicast Traffic in Network with Routers that run HSRP)
    - MSFC1
    - MSFC2
    - ARP and MAC Address Tables Before Ping Is Initiated
    - ARP and MAC Address Tables After Host A Sends ARP for Default Gateway
    - ARP and MAC Address Tables After Host B Sends ARP for Its Default Gateway
    - ARP and MAC Address Tables After Echo Packet Has Been Received by Host A
    - Consequences of Asymmetric Routing
    - ARP and MAC Address Tables After 5 Minutes of Continuous Ping of Host B by Host A
  - Case Study #6: HSRP Virtual IP Address Is Reported as a Different IP Address
  - Case Study #7: HSRP Causes MAC Violation on a Secure Port

Workaround
Case Study #8: %Interface Hardware Cannot Support Multiple Groups
Troubleshoot HSRP in Catalyst Switches
A. Verify HSRP Router Configuration
1. Verify Unique Router Interface IP Address
2. Verify Standby (HSRP) IP Addresses and Standby Group Numbers
3. Verify That Standby (HSRP) IP Address Is Different per Interface
4. When to Use the standy use-bia Command
5. Verify Access List Configuration
B. Verify Catalyst Fast EtherChannel and Trunking Configuration
1. Verify Trunking Configuration
2. Verify EtherChannel Configuration
3. Investigate Switch MAC Address Forwarding Table
C. Verify Physical Layer Connectivity
1. Check Interface Status
2. Link Change and Port Errors
3. Verify IP Connectivity
4. Check for Unidirectional Link
5. Additional Physical Layer Troubleshooting References
D. Layer 3 HSRP Debugging
1. Standard HSRP Debugging
2. Conditional HSRP Debugging (Limiting Output Based on Standby Group and/or VLAN)
3. Enhanced HSRP Debugging
E. Spanning Tree Troubleshooting
1. Verify Spanning Tree Configuration
2. Spanning Tree Loop Conditions
3. Topology Change Notification
4. Broadcast Suppression
5. Console and Telnet Access
6. Spanning Tree Features: Portfast, UplinkFast, and BackboneFast
7. BPDU Guard
8. VTP Pruning
F. Divide and Conquer
Related Information

Introduction

This document describes common issues and ways to troubleshoot Hot Standby Router Protocol (HSRP) problems.

Prerequisites

Requirements

There are no specific requirements for this document.

Components Used

This document includes examples from multiple Cisco IOS® Software releases and hardware platforms. Command availability and behavior vary by platform and release.

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.

Understand HSRP

Background Information

This document covers these most-common issues that relate to HSRP:

- Router report of a duplicate HSRP standby IP address
- Constant HSRP state changes (active, standby, speak)
- HSRP peers not present
- Switch error messages that relate to HSRP
- Excessive network unicast flooding to the HSRP configuration

---



Note: This document details how to troubleshoot HSRP in Catalyst switch environments. The document contains many references to software versions and network topology design. Nevertheless, the sole purpose of this document is to facilitate and guide engineers on how to troubleshoot HSRP. This document is not intended to be a design guide, software-recommendation document, or a best practices document.

Businesses and consumers that rely on intranet and Internet services for their mission-critical communications require and expect their networks and applications to be continuously available to them. Customers can satisfy their demands for near-100 percent network uptime if they leverage the HSRP in Cisco IOS® Software. HSRP, which is unique to Cisco platforms, provides network redundancy for IP networks in a manner that ensures that user traffic immediately and transparently recovers from first-hop failures in network edge devices or access circuits.

Two or more routers can act as a single, virtual router if they share an IP address and a MAC (Layer 2 [L2]) address. The address is necessary for host workstation default gateway redundancy. Most host workstations do not contain routing tables and use only a single next hop IP and MAC address. This address is known as a default gateway. With HSRP, members of the virtual router group continually exchange status messages.

One router can assume the routing responsibility of another if a router goes out of commission for either planned or unplanned reasons. Hosts are configured with a single default gateway and continue to forward IP packets to a consistent IP and MAC address. The changeover of devices that do the routing is transparent to the end workstations.

Note: You can configure host workstations that run Microsoft OS for multiple default gateways. But the multiple default gateways are not dynamic. The OS only uses one single default gateway at a time. The system only selects an additional configured default gateway at boot time if the first configured default gateway is determined unreachable by Internet Control Management Protocol (ICMP).

Basic Operation

A set of routers that run HSRP works in concert to present the illusion of a single default gateway router to the hosts on the LAN. This set of routers is known as an HSRP group or standby group. A single router that is elected from the group is responsible to forward 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. If the active router fails, the standby assumes the packet forwarding duties. Although an arbitrary number of routers can run HSRP, only the active router forwards the packets that are sent to the virtual router IP address.

In order to minimize network traffic, only the active and the standby routers send periodic HSRP messages after the protocol has completed the election process. Additional routers in the HSRP group remain in the `Listen` state. If the active router fails, the standby router takes over as the active router. If the standby router fails or becomes the active router, another router is elected as the standby router.

Each standby group emulates a single virtual router (default gateway). For each group, a single well-known MAC and IP address is allocated to that group. Multiple standby groups can coexist and overlap on a LAN, and individual routers can participate in multiple groups. In this case, the router maintains a separate state and timers for each group.

HSRP Terms

Term	Definition
Active router	The router that currently forwards packets for the virtual router
Standby router	The primary backup router
Standby group	The set of routers that participate in HSRP and jointly emulate a virtual router
Hello time	The interval between successive HSRP hello messages from a given router
Hold time	The interval between the receipt of a hello message and the presumption that the sending router has failed

HSRP Addressing

HSRP Router Communication

Routers that run HSRP communicate through HSRP packets. HSRPv1 uses IPv4 multicast address 224.0.0.2, and HSRPv2 uses IPv4 multicast address 224.0.0.102. Both versions use User Datagram Protocol (UDP) port 1985 (the addressing and examples in this document use HSRPv1 unless stated otherwise). IP multicast address 224.0.0.2 is a reserved multicast address that is used to communicate to all routers. The active router sources hello packets from its configured IP address and the HSRP virtual MAC address. The standby router sources hellos from its configured IP address and the burned-in MAC address (BIA). This use of source addressing is necessary so that HSRP routers can correctly identify each other.

In most cases, when you configure routers to be part of an HSRP group, the routers listen for the HSRP MAC address for that group as well as their own BIA. The only exception to this behavior is for Cisco 2500, 4000, and 4500 routers. These routers have Ethernet hardware that only recognizes a single MAC address. Therefore, these routers use the HSRP MAC address when they serve as the active router. The routers use their BIA when they serve as the standby router.

HSRP Standby IP Address Communication on All Media Except Token Ring

Because host workstations are configured with their default gateway as the HSRP standby IP address, hosts must communicate with the MAC address that is associated with the HSRP standby IP address. This MAC address is a virtual MAC address that is composed of 0000.0c07.ac**. The ** is the HSRP group number in hexadecimal, based on the respective interface. For example, HSRP group 1 uses the HSRP virtual MAC address of 0000.0c07.ac01. Hosts on the adjoining LAN segment use the normal Address Resolution Protocol (ARP) process in order to resolve the associated MAC addresses.

ICMP Redirects

HSRP peer routers that protect a subnet are able to provide access to all other subnets in the network. This is the basis of HSRP. Therefore, which router becomes the active HSRP router is irrelevant. In Cisco IOS software releases earlier than Cisco IOS Software Release 12.1(3)T, ICMP redirects are automatically disabled on an interface when HSRP is used on that interface. Without this configuration, the hosts can be redirected away from the HSRP virtual IP address and toward an interface IP and MAC address of a single router. Redundancy is lost.

Later Cisco IOS Software introduced a method to allow ICMP redirects with HSRP. This method filters outbound ICMP redirect messages through HSRP. The next hop IP address is changed to an HSRP virtual address. The gateway IP address in the outbound ICMP redirect message is compared to a list of HSRP active routers that are present on that network. If the router that corresponds to the gateway IP address is an active router for an HSRP group, the gateway IP address is replaced with that group virtual IP address. This solution allows hosts to learn optimal routes to remote networks and, at the same time, maintain the resilience that HSRP provides.

HSRP Functionality Matrix

Refer to the Cisco IOS® Release and HSRP Functionality Matrix section of Understand the Hot Standby Router Protocol Features and Functionality in order to learn about the features and Cisco IOS Software releases that support HSRP.

HSRP Features

This document provides information on these HSRP features:

- Preemption
- Interface tracking
- Use of a BIA
- Multiple HSRP groups
- Configurable MAC addresses
- Syslog support
- HSRP debugging
- Enhanced HSRP debugging
- Authentication
- IP redundancy
- Simple Network Management Protocol (SNMP) MIB
- HSRP for Multiprotocol Label Switching (MPLS)

Note: You can use your browser Find feature in order to locate these sections within the document.

Packet Format

This table shows the format of the data portion of the UDP HSRP frame:

Version	Op Code	State	Hellotime
Holdtime	Priority	Group	Reserved
Authentication Data			
Authentication Data			
Virtual IP Address			

This table describes each of the fields in the HSRP packet:

Packet Field	Description
Op Code (1 octet)	The Op Code describes the type of message that the packet contains. Possible values are: 0 - hello, 1 - coup, and 2 - resign. Hello messages are sent to indicate that a router runs HSRP and is able to become the active router. Coup messages are sent when a router wishes to become the active router. Resign messages are sent when a router no longer wishes to be the active router.
State (1 octet)	Each router in the standby group implements a state machine. The state field describes the current state of the router that sends the message. These are details on the individual

	states: 0 - initial, 1 - learn, 2 - listen, 4 - speak, 8 - standby, and 16 - active.
Hellotime (1 octet)	This field is only meaningful in hello messages. It contains the approximate period between the hello messages that the router sends. The time is given in seconds.
Holdtime (1 octet)	This field is only meaningful in hello messages. It contains the amount of time that the routers wait for a hello message before they initiate a state change.
Priority (1 octet)	This field is used to elect the active and standby routers. In a comparison of the priorities of two routers, the router with the highest value becomes the active router. The tie breaker is the router with the higher IP address.
Group (1 octet)	This field identifies the standby group.
Authentication Data (8 octets)	This field contains a cleartext, eight-character password.
Virtual IP Address (4 octets)	If the virtual IP address is not configured on a router, the address can be learned from the hello message from the active router. An address is only learned if no HSRP standby IP address has been configured, and the hello message is authenticated (if authentication is configured).

HSRP States

State	Definition
`Initial`	This is the state at the start. This state indicates that HSRP does not run. This state is entered through a configuration change or when an interface first becomes available.
`Learn`	The router has not determined the virtual IP address and has not yet seen an authenticated hello message from the active router. In this state, the router still waits to hear from the active router.
`Listen`	The router knows the virtual IP address, but the router is neither the active router nor the standby router. It listens for hello messages from those routers.
`Speak`	The router sends periodic hello messages and actively participates in the election of the active and/or standby router. A router cannot enter `speak` state unless the router has the virtual IP address.
`Standby`	The router is a candidate to become the next active router and sends periodic hello messages. With the exclusion of transient conditions, there is, at most, one router in the group in `standby` state.
`Active`	The router currently forwards packets that are sent to the group virtual MAC address. The router sends periodic hello messages. With the exclusion of transient conditions, there must be, at most, one router in `active` state in the group.

HSRP Timers

Each router uses three HSRP timers. The timers control hello transmission and peer expiration. HSRP convergence after a failure depends on how the HSRP hello and hold timers are configured. By default, these timers are set to 3 and 10 seconds, respectively, which means that a hello packet is sent between the HSRP standby group devices every 3 seconds, and the standby device becomes active when a hello packet has not been received for 10 seconds. You can lower these timer settings to speed up the failover or preemption, but, to avoid increased CPU usage and unnecessary standby state flapping, do not set the hello timer lower than one (1) second or the hold timer lower than 4 seconds.



Note: Use HSRP timer values supported by the platform and software release. Some newer platforms and releases support millisecond HSRP timers. Verify the supported timer range and assess CPU load, network conditions, and HSRP scale before deployment.

Note: If you use HSRP object tracking and the tracked object changes state, HSRP immediately applies the configured priority decrement or shutdown action. A role change occurs only when the resulting HSRP priorities, states, and preemption configuration trigger one. When a timer expires, the router transitions to a new HSRP state. The timers can be changed with this command: **standby <group-number> timers <hello-interval> <hold-time>**. For example, standby 1 timers 5 15.

This table provides more information on these timers:

Timer	Description
Active timer	This timer is used to monitor the active router. This timer starts any time an active router receives a hello packet. This timer expires in accordance with the hold time value that is set in the related field of the HSRP hello message.
Standby timer	This timer is used in order to monitor the standby router. The timer starts any time the standby router receives a hello packet. This timer expires in accordance with the hold time value that is set in the respective hello packet.
Hello timer	This timer is used to clock hello packets. HSRP routers in the Speak, Standby, and Active states generate a hello packet when the hello timer expires.

HSRP Events

This table provides the events in the HSRP finite state machine:

Key	Events
1	HSRP is configured on an enabled interface.
2	HSRP is disabled on an interface or the interface is disabled.
3	Active timer expiry The active timer is set to the hold time when the last hello message is seen from the active router.
4	Standby timer expiry The standby timer is set to the hold time when the last hello message is seen from the standby router.
5	Hello timer expiry The periodic timer for the send of hello messages is expired.
6	Receipt of a hello message of higher priority from a router in `speak` state
7	Receipt of a hello message of higher priority from the active router
8	Receipt of a hello message of lower priority from the active router
9	Receipt of a resign message from the active router
10	Receipt of a coup message from a higher priority router
11	Receipt of a hello message of higher priority from the standby router
12	Receipt of a hello message of lower priority from the standby router

HSRP Actions

This table specifies the actions to be taken as part of the state machine:

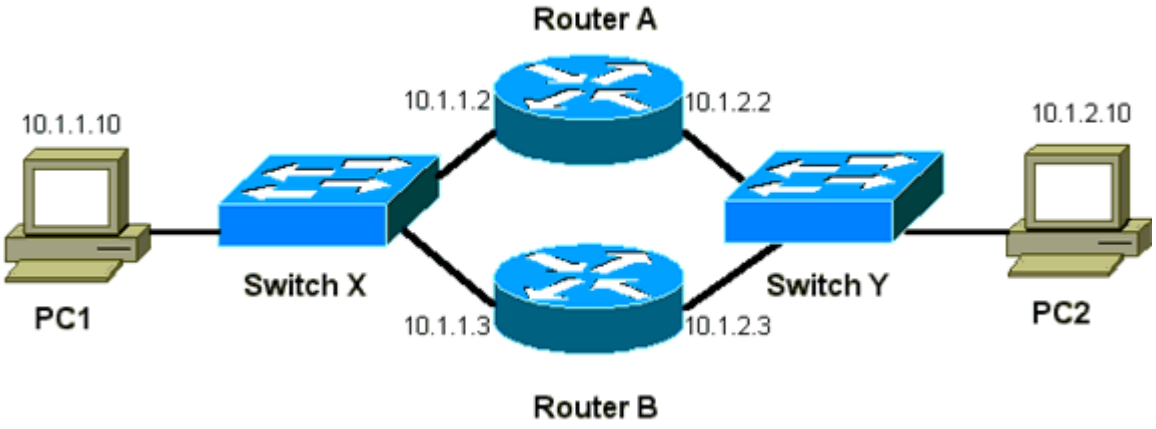
Letter	Action
A	Start active timer—If this action occurs as the result of the receipt of an authenticated hello message from the active router, the active timer is set to the hold time field in the hello message. Otherwise, the active timer is set to the current hold time value that is in use by this router. The active timer then starts.

B	Start standby timer—If this action occurs as the result of the receipt of an authenticated hello message from the standby router, the standby timer is set to the hold time field in the hello message. Otherwise, the standby timer is set to the current hold time value that is in use by this router. The standby timer then starts.
C	Stop active timer—The active timer stops.
D	Stop standby timer—The standby timer stops.
E	Learn parameters—This action is taken when an authenticated message is received from the active router. If the virtual IP address for this group is not manually configured, the virtual IP address can be learned from the message. The router can learn hello time and hold time values from the message.
F	Send hello message—The router sends a hello message with its current state, hello time, and hold time.
G	Send coup message—The router sends a coup message in order to inform the active router that there is a higher-priority router available.
H	Send resign message—The router sends a resign message in order to allow another router to become the active router.
I	Send gratuitous ARP message—The router broadcasts an ARP response packet that advertises the group virtual IP and MAC addresses. The packet is sent with the virtual MAC address as the source MAC address in the link layer header, as well as within the ARP packet.

HSRP State Table

The diagram in this section shows the state transitions of the HSRP state machine. Each time that an event occurs, the associated action results, and the router transitions to the next HSRP state. In the diagram, numbers designate events, and letters designate the associated action. The table in the section HSRP Events defines the numbers, and the table in the section HSRP Actions defines the letters. Use this diagram only as a reference. The diagram is detailed and is not necessary for general troubleshooting purposes.

Packet Flow



Device	MAC Address	IP Address	Subnet Mask	Default Gateway
PC1	0000.0c00.0001	10.1.1.10	255.255.255.0	10.1.1.1
PC2	0000.0c00.1110	10.1.2.10	255.255.255.0	10.1.2.1

Router A Configuration (Active Router)

```
interface GigabitEthernet 0/0
 ip address 10.1.1.2 255.255.255.0
 mac-address 4000.0000.0010
 standby 1 ip 10.1.1.1
 standby 1 priority 200

interface GigabitEthernet 0/1
 ip address 10.1.2.2 255.255.255.0
 mac-address 4000.0000.0011
 standby 1 ip 10.1.2.1
 standby 1 priority 200
```

Router B Configuration (Standby Router)

```
interface GigabitEthernet 0/0
    ip address 10.1.1.3 255.255.255.0
    mac-address 4000.0000.0020
    standby 1 ip 10.1.1.1

interface GigabitEthernet 0/1
     ip address 10.1.2.3 255.255.255.0
     mac-address 4000.0000.0021
     standby 1 ip 10.1.2.1
```

Note: These examples configure static MAC addresses for illustration purposes only. Do not configure static MAC addresses unless you are required to do so.

You must understand the concept behind packet flow when you obtain sniffer traces to troubleshoot HSRP problems. Router A uses the priority of 200 and becomes the active router on both interfaces. In the example in this section, packets from the router that are destined for a host workstation have the source MAC address of the router physical MAC address (BIA). Packets from the host machines that are destined for the HSRP IP address have the destination MAC address of the HSRP virtual MAC address. Note that the MAC addresses are not the same for each flow between the router and the host.

This table shows the respective MAC and IP address information per flow on the basis of a sniffer trace that is taken from Switch X.

Packet Flow	Source MAC	Destination MAC	Source IP	Destination IP
Packets from PC1 that are destined for PC2	PC1 (0000.0c00.0001)	HSRP virtual MAC address of Router A interface Ethernet 0 (0000.0c07.ac01)	10.1.1.10	10.1.2.10
Packets that return through Router A from PC2 and are destined for PC1	Router A Ethernet 0 BIA (4000.0000.0010)	PC1 (0000.0c00.0001)	10.1.2.10	10.1.1.10
Packets from PC1 that are destined for HSRP standby IP address (ICMP, Telnet)	PC1 (0000.0c00.0001)	HSRP virtual MAC address of Router A interface Ethernet 0 (0000.0c07.ac01)	10.1.1.10	10.1.1.1
Packets that are destined	PC1 (0000.0c00.0001)	Router A Ethernet 0 BIA	10.1.1.10	10.1.1.2

for the actual IP address of the active router (ICMP, Telnet)		(4000.0000.0010)		
Packets that are destined for the actual IP address of the standby router (ICMP, Telnet)	PC1 (0000.0c00.0001)	Router B Ethernet 0 BIA (4000.0000.0020)	10.1.1.10	10.1.1.3

Troubleshoot HSRP Case Studies

Case Study #1: HSRP Standby IP Address Is Reported as a Duplicate IP Address

These error messages can appear:

```
Oct 12 13:15:41: %STANDBY-3-DUPADDR: Duplicate address 10.25.0.1
  on Vlan25, sourced by 0000.0c07.ac19
Oct 13 16:25:41: %STANDBY-3-DUPADDR: Duplicate address 10.25.0.1
  on Vlan25, sourced by 0000.0c07.ac19
Oct 15 22:31:02: %STANDBY-3-DUPADDR: Duplicate address 10.25.0.1
  on Vlan25, sourced by 0000.0c07.ac19
Oct 15 22:41:01: %STANDBY-3-DUPADDR: Duplicate address 10.25.0.1
  on Vlan25, sourced by 0000.0c07.ac19
```

These error messages do not necessarily indicate an HSRP problem. Rather, the error messages indicate a possible Spanning Tree Protocol (STP) loop or router/switch configuration issue. The error messages are just symptoms of another problem.

In addition, these error messages do not prevent the proper operation of HSRP. The duplicate HSRP packet is ignored. These error messages are throttled at 30-second intervals. But, slow network performance and packet loss can result from the network instability that causes the `STANDBY-3-DUPADDR` error messages of the HSRP address.

These messages specifically indicate that the router received a data packet that was sourced from the HSRP IP address on VLAN 25 with the MAC addresses 0000.0c07.ac19. Since the HSRP MAC address is 0000.0c07.ac19, either the router in question received its own packet back or both routers in the HSRP group went into the `active` state. Because the router received its own packet, the problem most likely is with the network rather than the router. A variety of problems can cause this behavior. Among the possible network problems that cause the error messages are:

- Momentary STP loops
- EtherChannel configuration issues
- Duplicated frames

When you troubleshoot these error messages, see the steps to troubleshoot in the Troubleshooting HSRP in Catalyst Switches section of this document. All the troubleshoot modules are applicable to this section,

which includes modules on configuration. In addition, note any errors in the switch log and reference additional case studies as necessary.

You can use an access list in order to prevent the active router from receiving its own multicast hello packet. But, this is only a workaround for the error messages and actually hides the symptom of the problem. The workaround is to apply an extended inbound access list to the HSRP interfaces. The access list blocks all traffic that is sourced from the physical IP address and that is destined to all routers multicast address 224.0.0.2.

```
access-list 101 deny ip host 172.16.12.3 host 224.0.0.2
access-list 101 permit ip any any

interface GigabitEthernet 0/0
 ip address 172.16.12.3 255.255.255.0
 standby 1 ip 172.16.12.1
 ip access-group 101 in
```

Case Study #2: HSRP State Continuously Changes (Active, Standby, Speak) or %HSRP-6-STATECHANGE

These error messages can appear:

```
Jan 9 08:00:42.623: %STANDBY-6-STATECHANGE: Standby: 49:
  Vlan149 state Standby -> Active
Jan 9 08:00:56.011: %STANDBY-6-STATECHANGE: Standby: 49:
  Vlan149 state Active -> Speak
Jan 9 08:01:03.011: %STANDBY-6-STATECHANGE: Standby: 49:
  Vlan149 state Speak -> Standby
Jan 9 08:01:29.427: %STANDBY-6-STATECHANGE: Standby: 49:
  Vlan149 state Standby -> Active
Jan 9 08:01:36.808: %STANDBY-6-STATECHANGE: Standby: 49:
  Vlan149 state Active -> Speak
Jan 9 08:01:43.808: %STANDBY-6-STATECHANGE: Standby: 49:
  Vlan149 state Speak -> Standby
```

Jul 29 14:03:19.441: %HSRP-5-STATECHANGE: Vlan10 Grp 110 state Standby -> Active
Jul 29 16:27:04.133: %HSRP-5-STATECHANGE: Vlan10 Grp 110 state Active -> Speak
Jul 29 16:31:49.035: %HSRP-5-STATECHANGE: Vlan10 Grp 110 state Speak -> Standby

These error messages describe a situation in which a standby HSRP router did not receive HSRP hello packets from its peer before the active hold timer expired. The output shows that the standby router moves from the `standby` state to the `active` state. Shortly thereafter, the router returns to the `standby` state. Unless this error message occurs during the initial installation, an HSRP issue probably does not cause the error message. The error messages signify the loss of HSRP hellos between the peers. When you troubleshoot this issue, you must verify the communication between the HSRP peers. A random, momentary loss of data communication between the peers is the most common problem that results in these messages. HSRP state

changes are often due to High CPU Utilization. If the error message is due to high CPU utilization, put a sniffer on the network and the trace the system that causes the high CPU utilization.

There are several possible causes for the loss of HSRP packets between the peers. The most common problems are physical layer problems , excessive network traffic caused by spanning tree issues or excessive traffic caused by each Vlan. As with Case Study #1 , all the troubleshoot modules are applicable to the resolution of HSRP state changes, particularly the Layer 3 HSRP Debugging .

If the loss of HSRP packets between peers is due to excessive traffic caused by each VLAN as mentioned, you can tune or increase the SPD and hold the queue size to overcome the input queue drop problem.

To increase the Selective Packet Discard (SPD) size, go to the configuration mode and execute these commands on the Catalyst 6500 switches:

```
<#root>

(config)#

ip spd queue max-threshold 600


!--- Hidden Command

(config)#

ip spd queue min-threshold 500


!--- Hidden Command
```

In order to increase the hold queue size, go to the VLAN interface mode and execute this command:

```
<#root>

(config-if)#

hold-queue 500 in
```

After you increase the SPD and hold queue size, you can clear the interface counters if you execute the **clear counters <interface>** command.

Case Study #3: HSRP Does Not Recognize Peer

The router output in this section shows a router that is configured for HSRP but does not recognize its HSRP peers. In order for this to occur, the router must fail to receive HSRP hellos from the neighbor router. When

you troubleshoot this issue, see the Verify Physical Layer Connectivity section and the Verify HSRP Router Configuration section of this document. If the physical layer connectivity is correct, check for the mismatched VTP modes.

```
Vlan8 - Group 8
Local state is Active, priority 110, may preempt
Hellotime 3 holdtime 10
Next hello sent in 00:00:01.168
Hot standby IP address is 10.1.2.2 configured
Active router is local
Standby router is unknown expired
Standby virtual mac address is 0000.0c07.ac08
5 state changes, last state change 00:05:03
```

Case Study #4: HSRP State Changes and Switch Reports SYS-4-P2_WARN: 1/Host <mac_address> Is Flapping Between Port <port_1> and Port <port_2> in Syslog

These error messages can appear:

```
2001 Jan 03 14:18:43 %SYS-4-P2_WARN: 1/Host 00:00:0c:14:9d:08
  is flapping between port 2/4 and port 2/3
```

Feb 4 07:17:44 AST: %SW_MATM-4-MACFLAP_NOTIF: Host 0050.56a9.1f28 in vlan 1027 is flapping between port Te1/0/7 and port Te2/0/2

In Catalyst Switches, if the switch reports a host MAC address that moves between ports, and if the host MAC address moves twice within 15 seconds, a possible cause is an STP loop. The switch discards packets from this host for about 15 seconds in an effort to minimize the impact of an STP loop. If the reported MAC address that moves between two ports is the HSRP virtual MAC address, the problem is most likely an issue in which both HSRP routers go into the `active` state.

If the MAC address that is reported is not the HSRP virtual MAC address, the issue can indicate the loop, duplication, or reflection of packets in the network. These types of conditions can contribute to HSRP problems. The most common causes for the move of MAC addresses are spanning tree problems or physical layer problems .

When you troubleshoot this error message, complete these steps:

1. Determine the correct source (port) of the host MAC address.
2. Disconnect the port that must not source the host MAC address.
3. Document the STP topology on a per-VLAN basis and check for STP failure.
4. Verify the port channeling configuration.

1. An incorrect port channel configuration can result in the flap of error messages by the host MAC address. This is because of the load-balancing nature of port channeling.

Case Study #5: Asymmetric Routing and HSRP (Excessive Flooding of Unicast Traffic in Network with Routers that run HSRP)

With asymmetric routing, transmit and receive packets use different paths between a host and the peer with which it communicates. This packet flow is a result of the configuration of load-balancing between HSRP routers, based on HSRP priority, which set the HSRP to active or standby. This type of packet flow in a switching environment can result in excessive unknown unicast flooding. Also, Multilayer Switching (MLS) entries can be absent. Unknown unicast flooding occurs when the switch floods a unicast packet out of all ports. The switch floods the packet because there is no entry for the destination MAC address. This behavior does not break connectivity because packets are still forwarded. But, the behavior does account for the flood of extra packets on host ports. This case studies the behavior of asymmetric routing and why unicast flooding results.

Symptoms of asymmetric routing include:

- Excessive unicast packet flooding
- Absent MLS entry for flows
- Sniffer trace that shows that packets on the host port are not destined for the host
- Increased network latency with L2-based packet rewrite engines, such as server load balancers, web cache devices, and network appliances

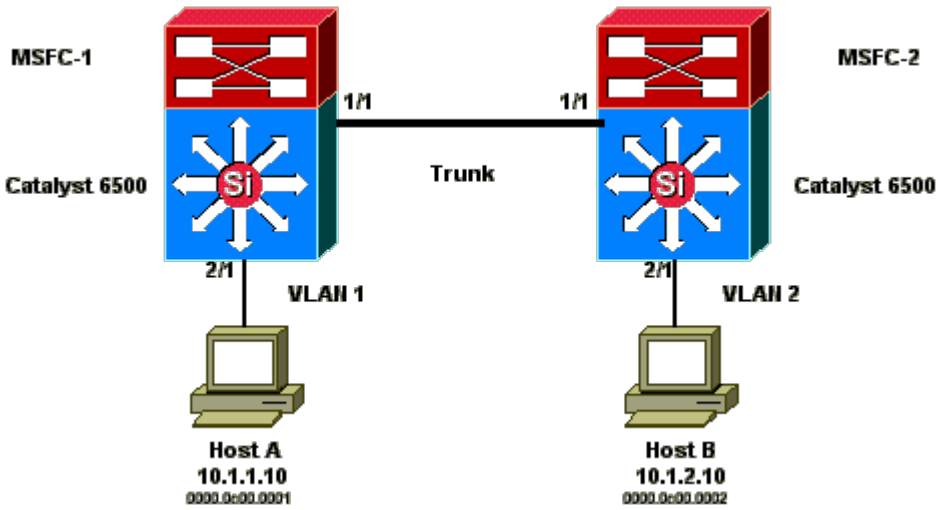
  Examples include the Cisco LocalDirector and Cisco Cache Engine.
- Dropped packets on connected hosts and workstations that cannot handle the additional unicast-flooding traffic load

---



Note: The default ARP cache aging time on a router is four hours. The default aging time of the switch content-addressable memory (CAM) entry is five minutes. The ARP aging time of the host workstations is not significant for this discussion. but, the example sets the ARP aging time to four hours.

This diagram illustrates this issue. This topology example includes Catalyst 6500s with Multilayer Switch Feature Cards (MSFCs) in each switch. Although this example uses MSFCs, you can use any router instead of the MSFC. Example routers that you can use include the Route Switch Module (RSM), Gigabit Switch Router (GSR), and Cisco 7500. The hosts are directly connected to ports on the switch. The switches are interconnected through a trunk which carries traffic for VLAN 1 and VLAN 2.



These outputs are excerpts from the **show standby** command output from each MSFC.

MSFC1

```
interface Vlan 1
 mac-address 0003.6bf1.2a01
 ip address 10.1.1.2 255.255.255.0
 no ip redirects
 standby 1 ip 10.1.1.1
 standby 1 priority 110

interface Vlan 2
 mac-address 0003.6bf1.2a01
 ip address 10.1.2.2 255.255.255.0
 no ip redirects
 standby 2 ip 10.1.2.1

MSFC1#show standby
Vlan1 - Group 1
Local state is Active, priority 110
Hellotime 3 holdtime 10
Next hello sent in 00:00:00.696
Hot standby IP address is 10.1.1.1 configured
Active router is local
Standby router is 10.1.1.3 expires in 00:00:07
Standby virtual mac address is 0000.0c07.ac01
2 state changes, last state change 00:20:40
Vlan2 - Group 2
Local state is Standby, priority 100
Hellotime 3 holdtime 10
Next hello sent in 00:00:00.776
Hot standby IP address is 10.1.2.1 configured
Active router is 10.1.2.3 expires in 00:00:09, priority 110
Standby router is local
4 state changes, last state change 00:00:51
MSFC1#exit
Console> (enable)
```

MSFC2

```
interface Vlan 1
 mac-address 0003.6bf1.2a02
 ip address 10.1.1.3 255.255.255.0
 no ip redirects
 standby 1 ip 10.1.1.1

interface Vlan 2
 mac-address 0003.6bf1.2a02
 ip address 10.1.2.3 255.255.255.0
 no ip redirects
 standby 2 ip 10.1.2.1
 standby 2 priority 110

MSFC2#show standby
Vlan1 - Group 1
Local state is Standby, priority 100
Hellotime 3 holdtime 10
Next hello sent in 00:00:01.242
Hot standby IP address is 10.1.1.1 configured
Active router is 10.1.1.2 expires in 00:00:09, priority 110
Standby router is local
7 state changes, last state change 00:01:17
Vlan2 - Group 2
Local state is Active, priority 110
Hellotime 3 holdtime 10
Next hello sent in 00:00:00.924
Hot standby IP address is 10.1.2.1 configured
Active router is local
Standby router is 10.1.2.2 expires in 00:00:09
Standby virtual mac address is 0000.0c07.ac02
2 state changes, last state change 00:40:08
MSFC2#exit
```



Note: On MSFC1, VLAN 1 is in the HSRP `active` state, and VLAN 2 is in the HSRP `standby` state. On MSFC2, VLAN 2 is in the HSRP `active` state, and VLAN 1 is in the HSRP `standby` state. The default gateway of each host is the respective standby IP address.

1. Initially, all caches are empty. Host A uses MSFC1 as its default gateway. Host B uses MSFC2.

ARP and MAC Address Tables Before Ping Is Initiated

Host A ARP Table	Switch 1 MAC Address Table MAC VLAN Port	MSFC1 ARP Table	MSFC2 ARP Table	Switch 2 MAC Address Table MAC VLAN Port	Host B ARP Table
	0003.6bf1.2a01 1 15/1			0003.6bf1.2a02 1 15/1	
	0003.6bf1.2a01 2 15/1			0003.6bf1.2a02 2 15/1	
	0000.0c07.ac01 1 15/1			0000.0c07.ac01 1 1/1	
	0000.0c07.ac02 2 1/1			0000.0c07.ac02 2 15/1	
	0003.6bf1.2a02 1 1/1			0003.6bf1.2a01 1 1/1	

	0003.6bf1.2a02 2 1/1			0003.6bf1.2a01 2 1/1	

Note: For brevity, the Switch 1 MAC address for the router HSRP and MAC address are not included in the other tables that appear in this section.

2. Host A pings host B, which means that host A sends an ICMP echo packet. Because each host resides on a separate VLAN, host A forwards its packets that are destined for host B to its default gateway. In order for that process to occur, host A must send an ARP in order to resolve its default gateway MAC address, 10.1.1.1.

ARP and MAC Address Tables After Host A Sends ARP for Default Gateway

Host A ARP Table	Switch 1 MAC Address Table MAC VLAN Port	MSFC1 ARP Table	MSFC2 ARP Table	Switch 2 MAC Address Table MAC VLAN Port	Host B ARP Table
10.1.1.1 : 0000.0c07.ac01	0000.0c00.0001 1 2/1	10.1.1.10 : 0000.0c00.0001			

3. MSFC1 receives the packet, rewrites the packet, and forwards the packet to host B. In order to rewrite the packet, MSFC1 sends an ARP request for host B because the host resides off a directly connected interface. MSFC2 has yet to receive any packets in this flow. When MSFC1 receives the ARP reply from host B, both switches learn the source port that is associated with host B.

ARP and MAC Address Tables After Host A Sends Packet to Default Gateway and MSFC1 Sends ARP for Host B

Host A ARP Table	Switch 1 MAC Address Table MAC VLAN Port	MSFC1 ARP Table	MSFC2 ARP Table	Switch 2 MAC Address Table MAC VLAN Port	Host B ARP Table
10.1.1.1 : 0000.0c07.ac01	0000.0c00.0001 1 2/1	10.1.1.10 : 0000.0c00.0001		0000.0c00.0002 2 2/1	10.1.2.2 : 0003.6bf1.2a01
	0000.0c00.0002 2 1/1	10.1.2.10 : 0000.0c00.0002			

4. Host B receives the echo packet from host A, through MSFC1. Host B must now send an echo reply to host A. Since host A resides on a different VLAN, host B forwards the reply through its default gateway, MSFC2. In order to forward the packet through, host B must send an ARP for its default gateway IP address, 10.1.2.1.

ARP and MAC Address Tables After Host B Sends ARP for Its Default Gateway

Host A ARP Table	Switch 1 MAC Address Table MAC VLAN Port	MSFC1 ARP Table	MSFC2 ARP Table	Switch 2 MAC Address Table MAC VLAN Port	Host B ARP Table
10.1.1.1 : 0000.0c07.ac01	0000.0c00.0001 1 2/1	10.1.1.10 : 0000.0c00.0001	10.1.2.10 0000.0c00.0002	0000.0c00.0002 2 2/1	10.1.2.2 (0003.6bf1.2a01)
	0000.0c00.0002 2 1/1	10.1.2.10 : 0000.0c00.0001			10.1.2.1 (0000.0c07.ac02)

5. Host B now forwards the echo reply packet to MSFC2. MSFC2 sends an ARP request for host A

because it is directly connected on VLAN 1. Switch 2 populates its MAC address table with the MAC address of host B.

ARP and MAC Address Tables After Echo Packet Has Been Received by Host A

Host A ARP Table	Switch 1 MAC Address Table MAC VLAN Port	MSFC1 ARP Table	MSFC2 ARP Table	Switch 2 MAC Address Table MAC VLAN Port	Host B ARP Table
10.1.1.1 : 0000.0c07.ac01	0000.0c00.0001 1 2/1	10.1.1.10 : 0000.0c00.0001	10.1.2.10 0000.0c00.0002	0000.0c00.0002 2 2/1	10.1.2.2 (0003.6bf1.2a01)
10.1.1.3 : 0003.6bf1.2a0	0000.0c00.0002 2 1/1	10.1.2.10 : 0000.0c00.0001	10.1.1.10 0000.0c00.0001	0000.0c00.00001 1 1/1	10.1.2.1 (0000.0c07.ac02)

6. The echo reply reaches host A and the flow is complete.

Consequences of Asymmetric Routing

Consider the case of the continuous ping of host B by host A. Remember that host A sends the echo packet to MSFC1, and host B sends the echo reply to MSFC2, which is in an asymmetric routing state. The only time that Switch 1 learns the source MAC of host B is when host B replies to an ARP request from MSFC1. This is because host B uses MSFC2 as its default gateway and does not send packets to MSFC1 and, consequently, Switch 1. Since the ARP timeout is four hours by default, Switch 1 ages the MAC address of host B after five minutes by default. Switch 2 ages host A after five minutes. As a result, Switch 1 must treat any packet with a destination MAC of host B as an unknown unicast. The switch floods the packet that comes from host A and is destined for host B out all ports. In addition, because there is no MAC address entry host B in Switch 1, there is no MLS entry as well.

ARP and MAC Address Tables After 5 Minutes of Continuous Ping of Host B by Host A

Host A ARP Table	Switch 1 MAC Address Table MAC VLAN Port	MSFC1 ARP Table	MSFC2 ARP Table	Switch 2 MAC Address Table MAC VLAN Port	Host B ARP Table
10.1.1.1 : 0000.0c07.ac01	0000.0c00.0001 1 2/1	10.1.1.10 : 0000.0c00.0001	10.1.2.10 0000.0c00.0002	0000.0c00.0002 2 2/1	10.1.2.2 : 0003.6bf1.2a01
10.1.1.3 : 0003.6bf1.2a0		10.1.2.10 : 0000.0c00.0001	10.1.1.10 0000.0c00.0001		10.1.2.1 : 0000.0c07.ac02

The echo reply packets that come from host B experience the same issue after the MAC address entry for host A ages on Switch 2. Host B forwards the echo reply to MSFC2, which in turn routes the packet and sends it out on VLAN 1. The switch does not have an entry host A in the MAC address table and must flood the packet out all ports in VLAN 1.

Asymmetric routing issues do not break connectivity. But, asymmetric routing can cause excessive unicast flooding and MLS entries that are missing. There are three configuration changes that can remedy this situation:

- Adjust the MAC aging time on the respective switches to 14,400 seconds (four hours) or longer.

- Change the ARP timeout on the routers to five minutes (300 seconds).
- Change the MAC aging time and ARP timeout to the same timeout value.

The preferable method is to change the MAC aging time to 14,400 seconds. These are the configuration guidelines:

- Cisco IOS Software:

  mac address-table aging-time <seconds> vlan <vlan_id>

Case Study #6: HSRP Virtual IP Address Is Reported as a Different IP Address

The "%STANDBY-3-DIFFVIP1" error message indicates that a router received an HSRP packet for the same group with a different virtual IP address. Check for inconsistent HSRP configuration, inter-VLAN leakage, and Layer 2 loops.

If you get this error message and there is interVLAN leakage because of bridging loops in the switch, complete these steps in order to resolve the error:

1. Identify the path that the packets take between end nodes.

   If there is a router on this path, complete these steps:

   1. Troubleshoot the path from the first switch to the router.
   2. Troubleshoot the path from the router to the second switch.

2. Connect to each switch on the path and check the status of the ports that are used on the path between end nodes.

Case Study #7: HSRP Causes MAC Violation on a Secure Port

When port security is configured on the switch ports that are connected to the HSRP enabled routers, it causes a MAC violation, since you cannot have the same secure MAC address on more than one interface. A security violation occurs on a secure port in one of these situations:

- The maximum number of secure MAC addresses is added to the address table, and a station whose MAC address is not in the address table attempts to access the interface.
- An address that is learned or configured on one secure interface is seen on another secure interface in the same VLAN.

By default, a port security violation causes the switch interface to become error-disabled and to shutdown immediately, which blocks the HSRP status messages between the routers.

Workaround

- Issue the **standby use-bia** command on the routers. This forces the routers to use a burned-in address for HSRP instead of the virtual MAC address.
- Disable port security on the switch ports that connect to the HSRP enabled routers.

Case Study #8: %Interface Hardware Cannot Support Multiple Groups

If multiple HSRP groups are created on the interface, this error message is received:

<#root>

```
%Interface hardware cannot support multiple groups
```

This error message is received due to the hardware limitation on some Routers or switches. It is not possible to overcome the limitation by any software methods. The problem is that each HSRP group uses one additional MAC address on interface, so the Ethernet MAC chip must support multiple programmable MAC addresses in order to enable several HSRP groups.

The workaround is to use the **standby use-bia** interface configuration command, which uses the Burned-In Address (BIA) of the interface as its virtual MAC address, instead of the preassigned MAC address.

Troubleshoot HSRP in Catalyst Switches

A. Verify HSRP Router Configuration

1. Verify Unique Router Interface IP Address

Verify that each HSRP router has a unique IP address for each subnet on a per-interface basis. Also, verify that each interface has the line protocol `up`. In order to quickly verify the current state of each interface, issue the **show ip interface brief** command. Here is an example:

<#root>

Router_1#

```
show ip interface brief
```

```
Interface              IP-Address      OK? Method Status                Protocol
Vlan1                  192.168.1.1     YES manual up                    up
```

```
Vlan10                 192.168.10.1    YES manual up                    up
Vlan11                 192.168.11.1    YES manual up                    up
```

Router_2#

show ip interface brief

```
Interface              IP-Address      OK? Method Status                Protocol
Vlan1                  192.168.1.2     YES manual up                    up
Vlan10                 192.168.10.2    YES manual up                    up
Vlan11                 192.168.11.2    YES manual up  up
```

2. Verify Standby (HSRP) IP Addresses and Standby Group Numbers

Verify that the configured standby (HSRP) IP addresses and standby group numbers match each HSRP-participating router. A mismatch of standby groups or HSRP standby addresses can cause HSRP problems. The **show standby** command details the standby group and standby IP address configuration of each interface. Here is an example:

<#root>

Router_1#

show standby

```
Vlan10 - Group 110
  State is Active
    2 state changes, last state change 00:01:34
  Virtual IP address is 192.168.10.100
  Active virtual MAC address is 0000.0c07.ac6e (MAC In Use)
    Local virtual MAC address is 0000.0c07.ac6e (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 0.144 secs
  Preemption enabled
  Active router is local
  Standby router is 192.168.10.2, priority 109 (expires in 10.784 sec)
  Priority 110 (configured 110)
  Group name is "hsrp-Vl10-110" (default)
  FLAGS: 0/1
Vlan11 - Group 111
  State is Active
    2 state changes, last state change 00:00:27
  Virtual IP address is 192.168.11.100
  Active virtual MAC address is 0000.0c07.ac6f (MAC In Use)
    Local virtual MAC address is 0000.0c07.ac6f (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 2.096 secs
  Preemption enabled
  Active router is local
  Standby router is 192.168.11.2, priority 109 (expires in 8.944 sec)
  Priority 110 (configured 110)
  Group name is "hsrp-Vl11-111" (default)
  FLAGS: 0/1
```

```
Router_2#

show standby


Vlan10 - Group 110
  State is Standby
    1 state change, last state change 00:03:15
  Virtual IP address is 192.168.10.100
  Active virtual MAC address is 0000.0c07.ac6e (MAC Not In Use)
    Local virtual MAC address is 0000.0c07.ac6e (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 1.088 secs
  Preemption disabled
  Active router is 192.168.10.1, priority 110 (expires in 11.584 sec)
  Standby router is local
  Priority 109 (configured 109)
  Group name is "hsrp-Vl10-110" (default)
  FLAGS: 0/1
Vlan11 - Group 111
  State is Standby
    1 state change, last state change 00:02:53
  Virtual IP address is 192.168.11.100
  Active virtual MAC address is 0000.0c07.ac6f (MAC Not In Use)
    Local virtual MAC address is 0000.0c07.ac6f (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 2.352 secs
  Preemption disabled
  Active router is 192.168.11.1, priority 110 (expires in 9.120 sec)
  Standby router is local
  Priority 109 (configured 109)
  Group name is "hsrp-Vl11-111" (default)
  FLAGS: 0/1
```

3. Verify That Standby (HSRP) IP Address Is Different per Interface

Verify that the standby (HSRP) IP address is unique from the configured IP address on each interface. The **show standby** command is a quick reference in order to view this information. Here is an example:

<#root>

Router_1#

show standby

```
Vlan10 - Group 110
  State is Active
    2 state changes, last state change 00:01:34
  Virtual IP address is 192.168.10.100
  Active virtual MAC address is 0000.0c07.ac6e (MAC In Use)
    Local virtual MAC address is 0000.0c07.ac6e (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 0.144 secs
```

```
  Preemption enabled
  Active router is local
  Standby router is 192.168.10.2, priority 109 (expires in 10.784 sec)
  Priority 110 (configured 110)
  Group name is "hsrp-Vl10-110" (default)
  FLAGS: 0/1
Vlan11 - Group 111
  State is Active
    2 state changes, last state change 00:00:27
  Virtual IP address is 192.168.11.100
  Active virtual MAC address is 0000.0c07.ac6f (MAC In Use)
    Local virtual MAC address is 0000.0c07.ac6f (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 2.096 secs
  Preemption enabled
  Active router is local
  Standby router is 192.168.11.2, priority 109 (expires in 8.944 sec)
  Priority 110 (configured 110)
  Group name is "hsrp-Vl11-111" (default)
  FLAGS: 0/1

Router_2#
```

show standby

```
Vlan10 - Group 110
  State is Standby
    1 state change, last state change 00:03:15
  Virtual IP address is 192.168.10.100
  Active virtual MAC address is 0000.0c07.ac6e (MAC Not In Use)
    Local virtual MAC address is 0000.0c07.ac6e (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 1.088 secs
  Preemption disabled
  Active router is 192.168.10.1, priority 110 (expires in 11.584 sec)
  Standby router is local
  Priority 109 (configured 109)
  Group name is "hsrp-Vl10-110" (default)
  FLAGS: 0/1
Vlan11 - Group 111
  State is Standby
    1 state change, last state change 00:02:53
  Virtual IP address is 192.168.11.100
  Active virtual MAC address is 0000.0c07.ac6f (MAC Not In Use)
    Local virtual MAC address is 0000.0c07.ac6f (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 2.352 secs
  Preemption disabled
  Active router is 192.168.11.1, priority 110 (expires in 9.120 sec)
  Standby router is local
  Priority 109 (configured 109)
  Group name is "hsrp-Vl11-111" (default)
  FLAGS: 0/1
```

4. When to Use the standy use-bia Command

Unless HSRP is configured on a Token Ring interface, only use the **standby use-bia** command in special circumstances. This command tells the router to use its BIA instead of the virtual HSRP MAC address for the HSRP group. On a Token Ring network, if source-route bridging (SRB) is in use, the **standby use-bia** command allows the new active router to update the host Routing Information Field (RIF) cache with a gratuitous ARP. But, not all of the host implementations handle the gratuitous ARP correctly. Another caveat for the **standby use-bia** command involves proxy ARP. A standby router cannot cover for the lost proxy ARP database of the failed active router.

5. Verify Access List Configuration

Verify that the access lists that are configured on all of the HSRP peers do not filter any HSRP addresses that are configured on their interfaces. Specifically, verify the multicast address that is used in order to send traffic to all of the routers on a subnet (**224.0.0.2**). Also, verify that the UDP traffic that is destined for the HSRP port **1985** is not filtered. HSRP uses this address and port to send hello packets between peers. Issue the **show access-lists** command as a quick reference to note the access lists that are configured on the router. Here is an example:

<#root>

```
Router_1#

show access-lists

Standard IP access list 77
    deny 10.19.0.0, wildcard bits 0.0.255.255
    permit any
Extended IP access list 144
    deny pim 238.0.10.0 0.0.0.255 any
    permit ip any any (58 matches)
```

B. Verify Catalyst Fast EtherChannel and Trunking Configuration

1. Verify Trunking Configuration

If HSRP peers communicate across one or more trunk links, verify that each trunk is operational and carries the HSRP VLAN. Run the show interfaces <interface> trunk command on each applicable switch and verify these conditions:

- The interface operates as an IEEE 802.1Q trunk.
- The native VLAN is consistent on both ends.
- The HSRP VLAN is allowed and active on the trunk.
- VTP or another VLAN-pruning mechanism does not remove the HSRP VLAN.
- STP places the trunk in the forwarding state for the HSRP VLAN.
- The switch does not report a native VLAN, encapsulation, or port-type inconsistency.

Configure both ends with compatible trunk settings. For a static switch-to-switch trunk, configure trunk mode explicitly on both ends. On platforms that support Dynamic Trunking Protocol (DTP), use dynamic trunk negotiation only when the network design requires it. The switchport nonegotiate command disables DTP negotiation; it does not independently configure the interface as a trunk.

<#root>

L2Switch_1#

show interfaces gigabitEthernet1/0/13 trunk

```
 Port Mode Encapsulation Status Native vlan Gi1/0/13 on 802.1q trunking 1 Port Vlans allowed on trunk G
```

Router_1#

show interfaces gigabitEthernet1/0/1 trunk

```
 Port Mode Encapsulation Status Native vlan Gi1/0/1 on 802.1q trunking 1 Port Vlans allowed on trunk Gi
```

2. Verify EtherChannel Configuration

If a port-channel is used in order to connect the HSRP routers, verify the EtherChannel configuration on both routers and switches. Configure both ends of an EtherChannel with compatible modes. Use desirable with desirable or auto for Port Aggregation Protocol (PAgP), active with active or passive for Link Aggregation Control Protocol (LACP), or on on both ends for a static EtherChannel.

However, in this example interfaces are not member of a port-channel:

<#root>

Router_1#

show etherchannel summary

```
Flags:  D - down        P - bundled in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        U - in use      f - failed to allocate aggregator

        M - not in use, minimum links not met
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port

        A - formed by Auto LAG


Number of channel-groups in use: 0
Number of aggregators:           0
```

```
Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------------------------------

Router_1#

Router_2#
```

show etherchannel summary

```
Flags:  D - down        P - bundled in port-channel
        I - stand-alone s - suspended
        H - Hot-standby (LACP only)
        R - Layer3      S - Layer2
        U - in use      f - failed to allocate aggregator

        M - not in use, minimum links not met
        u - unsuitable for bundling
        w - waiting to be aggregated
        d - default port

        A - formed by Auto LAG


Number of channel-groups in use: 0
Number of aggregators:           0

Group  Port-channel  Protocol    Ports
------+-------------+-----------+-----------------------------------------------

Router_2#
```

3. Investigate Switch MAC Address Forwarding Table

Verify that the MAC address table entries exist on the switch for the HSRP routers for the HSRP virtual MAC address and the physical BIAs. The **show standby** command on the router provides the virtual MAC address. The **show interface** command provides the physical BIA. Here are sample outputs:

<#root>

Router_1#

show standby

```
Vlan10 - Group 110
  State is Active
    2 state changes, last state change 00:37:03
  Virtual IP address is 192.168.10.100
  Active virtual MAC address is 0000.0c07.ac6e (MAC In Use)
    Local virtual MAC address is 0000.0c07.ac6e (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 0.768 secs
  Preemption enabled
  Active router is local
```

```
  Standby router is 192.168.10.2, priority 109 (expires in 10.368 sec)
  Priority 110 (configured 110)
  Group name is "hsrp-Vl10-110" (default)
  FLAGS: 0/1
Vlan11 - Group 111
  State is Active
    2 state changes, last state change 00:35:56
  Virtual IP address is 192.168.11.100
  Active virtual MAC address is 0000.0c07.ac6f (MAC In Use)
    Local virtual MAC address is 0000.0c07.ac6f (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 1.472 secs
  Preemption enabled
  Active router is local
  Standby router is 192.168.11.2, priority 109 (expires in 8.336 sec)
  Priority 110 (configured 110)
  Group name is "hsrp-Vl11-111" (default)
  FLAGS: 0/1
```

Router_1#

show interfaces vlan 10

```
Vlan10 is up, line protocol is up , Autostate Enabled
  Hardware is Ethernet SVI, address is d4e8.801f.4846 (bia d4e8.801f.4846)
  Internet address is 192.168.10.1/24
  MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive not supported
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:00, output 00:00:01, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/375/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     9258 packets input, 803066 bytes, 0 no buffer
     Received 0 broadcasts (0 IP multicasts)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     3034 packets output, 368908 bytes, 0 underruns
     Output 0 broadcasts (0 IP multicasts)
     0 output errors, 2 interface resets
     0 unknown protocol drops
     0 output buffer failures, 0 output buffers swapped out
```

L2Switch_1#

show mac address-table address 0000.0c07.ac6e

```
          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
  10    0000.0c07.ac6e    DYNAMIC     Gi1/0/13
Total Mac Addresses for this criterion: 1
```

```
L2Switch_1#

show mac address-table address 0000.0c07.ac6f

          Mac Address Table
-------------------------------------------

Vlan    Mac Address       Type        Ports
----    -----------       --------    -----
  11    0000.0c07.ac6f    DYNAMIC     Gi1/0/13
Total Mac Addresses for this criterion: 1
```

Be sure to check the CAM aging time in order to determine how quickly the entries are aged. If the time equals the configured value for STP forward delay, which is 15 seconds by default, there is a strong possibility that there is an STP loop in the network. Here is sample command output:

```
<#root>

L2Switch_1#

show mac address-table aging-time vlan 10

Global Aging Time:  300
Vlan    Aging Time
----    ----------
  10     300

L2Switch_1#

show mac address-table aging-time vlan 11

Global Aging Time:  300
Vlan    Aging Time
----    ----------
  11     300
```

C. Verify Physical Layer Connectivity

If more than one router in an HSRP group becomes active, those routers do not consistently receive the hello packets from fellow HSRP peers. Physical layer problems can prevent the consistent pass of traffic between peers and cause this scenario. Be sure to verify physical connectivity and IP connectivity between HSRP peers when you troubleshoot HSRP. Issue the **show standby** command in order to verify connectivity. Here is an example:

```
<#root>

Router_1#

show standby
```

```
Vlan10 - Group 110
  State is Active
    2 state changes, last state change 00:54:03
  Virtual IP address is 192.168.10.100
  Active virtual MAC address is 0000.0c07.ac6e (MAC In Use)
    Local virtual MAC address is 0000.0c07.ac6e (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 0.848 secs
  Preemption enabled
  Active router is local
  Standby router is unknown
  Priority 110 (configured 110)
  Group name is "hsrp-Vl10-110" (default)
  FLAGS: 0/1
Vlan11 - Group 111
  State is Active
    2 state changes, last state change 00:52:56
  Virtual IP address is 192.168.11.100
  Active virtual MAC address is 0000.0c07.ac6f (MAC In Use)
    Local virtual MAC address is 0000.0c07.ac6f (v1 default)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 0.512 secs
  Preemption enabled
  Active router is local
  Standby router is unknown
  Priority 110 (configured 110)
  Group name is "hsrp-Vl11-111" (default)
  FLAGS: 0/1
```

Router_2#

show standby

```
Vlan10 - Group 110
  State is Init (interface down)
    2 state changes, last state change 00:00:42
  Virtual IP address is 192.168.10.100
  Active virtual MAC address is unknown (MAC Not In Use)
    Local virtual MAC address is 0000.0c07.ac6e (v1 default)
  Hello time 3 sec, hold time 10 sec
  Preemption disabled
  Active router is unknown
  Standby router is unknown
  Priority 109 (configured 109)
  Group name is "hsrp-Vl10-110" (default)
  FLAGS: 0/1
Vlan11 - Group 111
  State is Init (interface down)
    2 state changes, last state change 00:00:36
  Virtual IP address is 192.168.11.100
  Active virtual MAC address is unknown (MAC Not In Use)
    Local virtual MAC address is 0000.0c07.ac6f (v1 default)
  Hello time 3 sec, hold time 10 sec
  Preemption disabled
  Active router is unknown
  Standby router is unknown
  Priority 109 (configured 109)
  Group name is "hsrp-Vl11-111" (default)
  FLAGS: 0/1
```

1. Check Interface Status

Check the interfaces. Verify that all HSRP-configured interfaces are `up/up`, as this example shows:

```
<#root>

Router_1#

show ip interface brief


Interface              IP-Address      OK? Method Status                Protocol
Vlan1                  192.168.1.1     YES manual up                    up
Vlan10                 192.168.10.1    YES manual up                    up
Vlan11                 192.168.11.1    YES manual up                    up

Router_2#

show ip interface brief


Interface              IP-Address      OK? Method Status                Protocol
Vlan1                  192.168.1.2     YES manual up                    up
Vlan10                 192.168.10.2    YES manual administratively down down
Vlan11                 192.168.11.2    YES manual administratively down down
```

If any interfaces are administratively `down/down`, enter the configuration mode on the router and issue the **no shutdown** interface-specific command. Here is an example:

```
<#root>

Router_2#

configure terminal


Enter configuration commands, one per line.  End with CNTL/Z.
Router_2(config)#

interface vlan 10


Router_2(config-if)#

no shutdown


Router_2(config-if)#
```

```
end
```

```
Router_2#
```

```
configure terminal
```

```
Enter configuration commands, one per line.  End with CNTL/Z.
Router_2(config)#
```

```
interface vlan 11
```

```
Router_2(config-if)#
```

```
no shutdown
```

```
Router_2(config-if)#
```

```
end
```

```
Router_2#
```

```
show ip interface brief
```

```
Interface              IP-Address      OK? Method Status                Protocol
Vlan1                  192.168.1.2     YES manual up                    up
Vlan10                 192.168.10.2    YES manual up                    down
Vlan11                 192.168.11.2    YES manual up                    up
```

If any interfaces are `down/down` or `up/down`, review the log for any interface change notifications. For Cisco IOS Software-based switches, these messages appear for link `up/down` situations:

<#root>

```
%LINK-3-UPDOWN: Interface "interface", changed state to up
%LINK-3-UPDOWN: Interface "interface", changed state to down
```

```
Router_1#
```

```
show logging
```

```
3d04h: %STANDBY-6-STATECHANGE: Standby: 0: Vlan10 state Active-> Speak
3d04h: %LINK-5-CHANGED: Interface Vlan10, changed state to down
3d04h: %LINEPROTO-5-UPDOWN: Line protocol on Interface Vlan10, changed state to down
```

Inspect the ports, cables, and any transceivers or other devices that are between the HSRP peers. Has anyone removed or loosened any connections? Are there any interfaces that lose a link repeatedly? Are the proper cable types used? Check the interfaces for any errors, as this example shows:

<#root>

Router_2#

show interface vlan 10

Vlan10 is down, line protocol is down , Autostate Enabled
  Hardware is Ethernet SVI, address is 1880.90d8.5946 (bia 1880.90d8.5946)
  Internet address is 192.168.10.2/24
  MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive not supported
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:10, output 00:00:08, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/375/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 0 bits/sec, 0 packets/sec
  5 minute output rate 0 bits/sec, 0 packets/sec
     1243 packets input, 87214 bytes, 0 no buffer
     Received 0 broadcasts (0 IP multicasts)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     23 packets output, 1628 bytes, 0 underruns
     Output 0 broadcasts (0 IP multicasts)
     0 output errors, 2 interface resets
     0 unknown protocol drops
     0 output buffer failures, 0 output buffers swapped out

2. Link Change and Port Errors

Check the switch ports link changes and other errors. Issue these commands and review the output:

- **show logging**
- **show interfaces <interface> counters**
- **show interfaces <interface> status**

These commands help you determine if there is a problem with connectivity between switches and other devices.

These messages are normal for link `up/down` situations:

<#root>

L2Switch_1#

show logging

```
Syslog logging: enabled (0 messages dropped, 5 messages rate-limited, 0 flushes, 0 overruns, xml disabl

No Active Message Discriminator.



No Inactive Message Discriminator.


    Console logging: level informational, 319 messages logged, xml disabled,
                     filtering disabled
    Monitor logging: level debugging, 0 messages logged, xml disabled,
                     filtering disabled
    Buffer logging:  level debugging, 467 messages logged, xml disabled,
                    filtering disabled
    Exception Logging: size (4096 bytes)
    Count and timestamp logging messages: disabled
    File logging: disabled
    Persistent logging: disabled

No active filter modules.

    Trap logging: level informational, 327 message lines logged
        Logging Source-Interface:       VRF Name:

Log Buffer (10000 bytes):
*Jul 26 17:52:07.526: %LINK-3-UPDOWN: Interface GigabitEthernet1/0/13, changed state to up
*Jul 26 17:52:09.747: %LINK-3-UPDOWN: Interface GigabitEthernet1/0/13, changed state to down
*Jul 26 17:57:11.716: %SPANTREE-7-RECV_1Q_NON_TRUNK: Received 802.1Q BPDU on non trunk GigabitEthernet1
*Jul 26 17:57:11.716: %SPANTREE-7-BLOCK_PORT_TYPE: Blocking GigabitEthernet1/0/16 on VLAN0307. Inconsis
*Jul 26 17:57:13.583: %LINK-3-UPDOWN: Interface GigabitEthernet1/0/16, changed state to up
*Jul 26 17:57:16.237: %LINK-3-UPDOWN: Interface GigabitEthernet1/0/16, changed state to down
*Jul 26 18:02:16.481: %SPANTREE-7-RECV_1Q_NON_TRUNK: Received 802.1Q BPDU on non trunk GigabitEthernet1
*Jul 26 18:02:16.481: %SPANTREE-7-BLOCK_PORT_TYPE: Blocking GigabitEthernet1/0/16 on VLAN0307. Inconsis
*Jul 26 18:02:18.367: %LINK-3-UPDOWN: Interface GigabitEthernet1/0/16, changed state to up
*Jul 26 18:02:20.561: %LINK-3-UPDOWN: Interface GigabitEthernet1/0/16, changed state to down
```

Run the **show interfaces <interface> status** command in order to determine the general health of a port. Here is an example:

<#root>

L2Switch_1#

show interfaces gigabitEthernet 1/0/13 status

```
Port      Name               Status       Vlan       Duplex  Speed Type
Gi1/0/13                     connected    trunk      a-full a-1000 10/100/1000BaseTX
```

Is the interface status `connected`, `notconnect`, or `errdisable`? If the status is `notconnect`, check that the cable is plugged in on both sides. Check that the proper cable is used. If the status is `errdisable`, review the counters for excessive errors. Refer to Recover Errdisable Port State on Cisco IOS Platforms for more information.

For what VLAN is this port configured? Be sure that the other side of the connection is configured for the same VLAN. If the link is configured to be a trunk, be sure that both sides of the trunk carry the same VLANs.

What is the speed and duplex configuration? If the setting is preceded by a-, the port is configured to autonegotiate the speed and duplex. Otherwise, the network administrator has predetermined this configuration. For configuration of the speed and duplex for a link, the settings on both sides of the link must match. If one switch port is configured for autonegotiation, the other side of the link must also be configured for autonegotiation. If one side is hard coded to a specific speed and duplex, the other side must be hard coded as well. If you leave one side to autonegotiate while the other side is hard coded, you break the autonegotiation process.

<#root>

L2Switch_1#

show interfaces gi1/0/13 counters errors

```
 Port Align-Err FCS-Err Xmit-Err Rcv-Err UnderSize OutDiscards Gi1/0/13
 0 0 0 0 0 0

 Port Single-Col Multi-Col Late-Col Excess-Col Carri-Sen Runs Gi1/0/13
0 0 0 0 0 0
```

Are there many `Align-Err`, `FCS-Err`, or `Runts`? These can indicate a speed or duplex mismatch, cabling fault, transceiver fault, or another physical-layer problem. Change the speed and duplex settings for that port in order to help correct these errors.

Run the **show interfaces <interface> counters** command in order to verify that the port is passing traffic. The `In` and `Out` columns indicate the number of unicast, multicast, and broadcast packets that are received and transmitted on a particular port. The bottom counters reveal how many packets are discarded or lost and whether these packets are a part of inbound or outbound traffic. **`Lrn-Discrd`, `In-Lost`**, and **`Out-Lost`** count the number of packets that are mistakenly forwarded or dropped due to insufficient buffers.

<#root>

L2Switch_1#

```
show interfaces gi1/0/13 counters



Port            InOctets    InUcastPkts    InMcastPkts    InBcastPkts
Gi1/0/13       304933333        1180453        1082538          14978

Port           OutOctets   OutUcastPkts   OutMcastPkts   OutBcastPkts
Gi1/0/13       282752538         276716         824562         588960
```

3. Verify IP Connectivity

Verify IP connectivity. Issue an IP ping from the associated router to the remote HSRP device. This helps expose any momentary losses of connectivity. An extended ping is only available in enable mode. Here is sample command output:

<#root>

Router_1#

```
show run interface vlan 10


Building configuration...

Current configuration : 141 bytes
!
interface Vlan10
 ip address 192.168.10.1 255.255.255.0
 standby 110 ip 192.168.10.100
 standby 110 priority 110
 standby 110 preempt
end

Router_2#

show run interface vlan 10


Building configuration...

Current configuration : 120 bytes
!
interface Vlan10
 ip address 192.168.10.2 255.255.255.0
 standby 110 ip 192.168.10.100
 standby 110 priority 109
end

Router_1#

ping 192.168.10.2 repeat 1500


Type escape sequence to abort.
Sending 1500, 100-byte ICMP Echos to 192.168.10.2, timeout is 2 seconds:
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
```

```
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
Success rate is 100 percent (1500/1500), round-trip min/avg/max = 1/2/9 ms
```

Issue the ping from each HSRP router to its peers in order to determine the location of the failure in connectivity.

4. Check for Unidirectional Link

Check the switch for unidirectional links between HSRP peers. A unidirectional link occurs whenever traffic that a local device transmits over a link is received by the neighbor, but traffic that the neighbor transmits is not received by the local device. This feature is known as UniDirectional Link Detection (UDLD) aggressive mode. Use of UDLD is only possible if both sides of the connection support the feature. UDLD aggressive mode operates at L2 to determine if a link is connected correctly and if traffic flows bidirectionally between the right neighbors. Here are sample command outputs:

Note: Navigate to the next link to Configure the UDLD Protocol Feature It depends on which platform is used.

Another option that can help to verify a Unidirectional Link if UDLD is not available is with the use of Cisco Discovery Protocol (CDP). Enablement of CDP is another way to detect if a unidirectional link exists. If only one side of a link can see its neighbor device, replace the cable between the devices and check for faulty interfaces.

<#root>

Router_1#

show cdp

Global CDP information:
    Sending CDP packets every 60 seconds

```
    Sending a holdtime value of 180 seconds
    Sending CDPv2 advertisements is  enabled

Router_1#show cdp neighbors gi1/0/1 detail
-------------------------
Device ID: L2Switch_1.cisco.com
Entry address(es):
  IP address: 192.168.70.1
  IPv6 address: 2001:420:140E:2101::1  (global unicast)
  IPv6 address: FE80::2FE:C8FF:FED3:86C7  (link-local)
Platform: cisco WS-C3650-12X48UR,  Capabilities: Router Switch IGMP
Interface: GigabitEthernet1/0/1,  Port ID (outgoing port): GigabitEthernet1/0/13
Holdtime : 173 sec

Version :
Cisco IOS Software [Denali], Catalyst L3 Switch Software (CAT3K_CAA-UNIVERSALK9-M), Version 16.3.8, RELEASE SOFTWARE (fc3)
Technical Support: http://www.cisco.com/techsupport
Copyright (c) 1986-2019 by Cisco Systems, Inc.
Compiled Wed 13-Feb-19 03:00 by mcpre

advertisement version: 2
VTP Management Domain: 'CALOnet'
Native VLAN: 1
Duplex: full
Management address(es):
  IP address: 192.168.70.1
Spare Pair PoE: Yes, Spare Pair Detection Required: No
Spare Pair PD Config: Disable, Spare Pair PSE Operational: No


Total cdp entries displayed : 1
```

5. Additional Physical Layer Troubleshooting References

Refer to these documents:

- Configure and Verify Ethernet 10/100/1000Mb Half/Full Duplex Auto-Negotiation
- Recover Errdisable Port State on Cisco IOS Platforms
- The Understanding Data Link Errors section of Troubleshoot Catalyst Switches to NIC Compatibility Issues
- Troubleshoot Switch Port and Interface Problems

D. Layer 3 HSRP Debugging

If the HSRP state changes are frequent, use the HSRP debug commands (in enable mode) on the router in order to watch HSRP activity. This information helps you determine what HSRP packets are received and sent by the router. Gather this information if you create a service request with Cisco Technical Support. The debug output also shows HSRP state information, along with detailed HSRP hello packet accounts.

1. Standard HSRP Debugging

In Cisco IOS enable the HSRP debug capability with the command **debug standby**. This information is useful where problems are intermittent and affect only a few interfaces. The debug enables you to determine if the HSRP router in question receives and transmits HSRP hello packets at specific intervals. If the router does not receive hello packets, you can infer that either the peer does not transmit the hello packets or the network drops the packets.

Command	Purpose
debug standby	Enables HSRP debugging

Here is sample command output:

<#root>

Router_1#

debug standby

```
HSRP debugging is on
Jul 29 16:12:16.889: HSRP: Vl10 Grp 110 Hello  out 192.168.10.1 Active  pri 110 vIP 192.168.10.100
Jul 29 16:12:16.996: HSRP: Vl11 Grp 111 Hello  in  192.168.11.2 Standby pri 109 vIP 192.168.11.100
Jul 29 16:12:17.183: HSRP: Vl10 Grp 110 Hello  in  192.168.10.2 Standby pri 109 vIP 192.168.10.100
Jul 29 16:12:17.366: HSRP: Vl11 Grp 111 Hello  out 192.168.11.1 Active  pri 110 vIP 192.168.11.100
Jul 29 16:12:18.736: HSRP: Vl10 Interface adv in, Passive, active 0, passive 1, from 192.168.10.2
Jul 29 16:12:19.622: HSRP: Vl10 Grp 110 Hello  out 192.168.10.1 Active  pri 110 vIP 192.168.10.100
```

2. Conditional HSRP Debugging (Limiting Output Based on Standby Group and/or VLAN)

Cisco IOS Software Release 12.0(3) introduced a debug condition to allow the output from the **debug standby** command to be filtered based on interface and group number. The command utilizes the debug condition paradigm that was introduced in Cisco IOS Software Release 12.0.

Command	Purpose
debug condition standby <interface> <group>	Enables HSRP conditional debugging of the group (0–255)

The interface must be a valid interface that can support HSRP. The group can be any group, from 0 through 255. A debug condition can be set for groups that do not exist. This allows debugs to be captured during the initialization of a new group. Debug standby must be enabled in order to produce any debug output. If no standby debug conditions exist, debug output is produced for all groups on all interfaces. If at least one standby debug condition exists, standby debug output is filtered based on all of the standby debug conditions. Here is sample command output:

```
Router_1#debug condition standby vlan 10 110
Condition 1 set
Router_1#
```

```
Jul 29 16:16:20.284: Vl10 HSRP110 Debug: Condition 1, hsrp Vl10 HSRP110 triggered, count 1
Router_1#debug standby
HSRP debugging is on
Router_1#
Jul 29 16:16:44.797: HSRP: Vl10 Grp 110 Hello out 192.168.10.1 Active pri 110 vIP 192.168.10.100
Jul 29 16:16:45.381: HSRP: Vl10 Grp 110 Hello in 192.168.10.2 Standby pri 109 vIP 192.168.10.100
Jul 29 16:16:47.231: HSRP: Vl10 Grp 110 Hello out 192.168.10.1 Active pri 110 vIP 192.168.10.100
Jul 29 16:16:48.248: HSRP: Vl10 Grp 110 Hello in 192.168.10.2 Standby pri 109 vIP 192.168.10.100
```

3. Enhanced HSRP Debugging

Cisco IOS Software Release 12.1(1) added enhanced HSRP debugging. In order to help find useful information, enhanced HSRP debugging limits the noise from periodic hello messages and includes additional state information. This information is particularly useful when you work with a Cisco Technical Support engineer if you create a service request.

Command	Purpose
debug standby	Displays all HSRP errors, events, and packets
debug standby errors	Displays HSRP errors
debug standby events [[all] | [hsrp | redundancy | track]] [detail]	Displays HSRP events
debug standby packets [[all | terse] | [advertise | coup | hello | resign]] [detail]	Displays HSRP packets
debug standby terse	Display limited range of HSRP errors, events and packets

Here is sample command output:

<#root>

Router_2#

debug standby terse

```
HSRP:
  HSRP Errors debugging is on
  HSRP Events debugging is on
    (protocol, neighbor, redundancy, track, ha, arp, interface)
  HSRP Packets debugging is on
    (Coup, Resign)
Router_2#
*Jul 29 16:49:35.416: HSRP: Vl10 Grp 110 Resign in  192.168.10.1 Active  pri 110 vIP 192.168.10.100
*Jul 29 16:49:35.416: HSRP: Vl10 Grp 110 Standby: i/Resign rcvd (110/192.168.10.1)
*Jul 29 16:49:35.416: HSRP: Vl10 Grp 110 Active router is local, was 192.168.10.1
*Jul 29 16:49:35.416: HSRP: Vl10 Nbr 192.168.10.1 no longer active for group 110 (Standby)
*Jul 29 16:49:35.417: HSRP: Vl10 Nbr 192.168.10.1 Was active or standby - start passive holddown
*Jul 29 16:49:35.417: HSRP: Vl10 Grp 110 Standby router is unknown, was local
*Jul 29 16:49:35.417: HSRP: Vl10 Grp 110 Standby -> Active
*Jul 29 16:49:35.418: %HSRP-5-STATECHANGE: Vlan10 Grp 110 state Standby -> Active
*Jul 29 16:49:35.418: HSRP: Peer not present
*Jul 29 16:49:35.418: HSRP: Vl10 Grp 110 Redundancy "hsrp-Vl10-110" state Standby -> Active
```

```
*Jul 29 16:49:35.419: HSRP: Vl10 Grp 110 Added 192.168.10.100 to ARP (0000.0c07.ac6e)
*Jul 29 16:49:35.420: HSRP: Vl10 IP Redundancy "hsrp-Vl10-110" standby, local -> unknown
*Jul 29 16:49:35.421: HSRP: Vl10 IP Redundancy "hsrp-Vl10-110" update, Standby -> Active
*Jul 29 16:49:38.422: HSRP: Vl10 IP Redundancy "hsrp-Vl10-110" update, Active -> Active
```

You can use interface and/or HSRP group conditional debugging in order to filter this debug output.

Command	Purpose
debug condition interface interface	Enables interface conditional debugging
debug condition standby <interface> <group>	Enables HSRP conditional debugging

In this example, the router joins a preexisting HSRP group:

<#root>

Rotuer_2#

debug condition standby vlan 10 110

```
Condition 1 set
Router_2#
```

debug condition interface gigabitEthernet 1/0/1 vlan-id 10

```
Condition 2 set
Router_2#
```

debug standby

```
HSRP debugging is on
Router_2#
*Jul 29 16:54:12.496: HSRP: Vl10 Grp 110 Hello  out 192.168.10.2 Active  pri 109 vIP 192.168.10.100
*Jul 29 16:54:15.122: HSRP: Vl10 Grp 110 Hello  out 192.168.10.2 Active  pri 109 vIP 192.168.10.100
*Jul 29 16:54:17.737: HSRP: Vl10 Grp 110 Hello  out 192.168.10.2 Active  pri 109 vIP 192.168.10.100
*Jul 29 16:54:18.880: HSRP: Vl10 Nbr 192.168.10.1 is passive
*Jul 29 16:54:20.316: HSRP: Vl10 Grp 110 Hello  out 192.168.10.2 Active  pri 109 vIP 192.168.10.100
*Jul 29 16:54:20.322: HSRP: Vl10 Grp 110 Coup   in  192.168.10.1 Listen  pri 110 vIP 192.168.10.100
*Jul 29 16:54:20.323: HSRP: Vl10 Grp 110 Active: j/Coup rcvd from higher pri router (110/192.168.10.1)
*Jul 29 16:54:20.323: HSRP: Vl10 Grp 110 Active router is 192.168.10.1, was local
*Jul 29 16:54:20.323: HSRP: Vl10 Nbr 192.168.10.1 is no longer passive
*Jul 29 16:54:20.324: HSRP: Vl10 Nbr 192.168.10.1 active for group 110
*Jul 29 16:54:20.324: HSRP: Vl10 Grp 110 Active -> Speak
*Jul 29 16:54:20.325: %HSRP-5-STATECHANGE: Vlan10 Grp 110 state Active -> Speak
*Jul 29 16:54:20.325: HSRP: Peer not present
*Jul 29 16:54:20.325: HSRP: Vl10 Grp 110 Redundancy "hsrp-Vl10-110" state Active -> Speak
*Jul 29 16:54:20.326: HSRP: Vl10 Grp 110 Removed 192.168.10.100 from ARP
*Jul 29 16:54:20.326: HSRP: Vl10 Grp 110 Deactivating MAC 0000.0c07.ac6e
*Jul 29 16:54:20.327: HSRP: Vl10 Grp 110 Removing 0000.0c07.ac6e from MAC address filter
*Jul 29 16:54:20.328: HSRP: Vl10 Grp 110 Hello  out 192.168.10.2 Speak   pri 109 vIP 192.168.10.100
*Jul 29 16:54:20.328: HSRP: Vl10 Grp 110 Hello  in  192.168.10.1 Active  pri 110 vIP 192.168.10.100
*Jul 29 16:54:23.104: HSRP: Vl10 Grp 110 Hello  out 192.168.10.2 Speak   pri 109 vIP 192.168.10.100
*Jul 29 16:54:23.226: HSRP: Vl10 Grp 110 Hello  in  192.168.10.1 Active  pri 110 vIP 192.168.10.100
*Jul 29 16:54:25.825: HSRP: Vl10 Grp 110 Hello  in  192.168.10.1 Active  pri 110 vIP 192.168.10.100
```

```
*Jul 29 16:54:25.952: HSRP: Vl10 Grp 110 Hello  out 192.168.10.2 Speak   pri 109 vIP 192.168.10.100
*Jul 29 16:54:28.427: HSRP: Vl10 Grp 110 Hello  in  192.168.10.1 Active  pri 110 vIP 192.168.10.100
*Jul 29 16:54:28.772: HSRP: Vl10 Grp 110 Hello  out 192.168.10.2 Speak   pri 109 vIP 192.168.10.100
*Jul 29 16:54:30.727: HSRP: Vl10 Grp 110 Speak: d/Standby timer expired (unknown)
*Jul 29 16:54:30.727: HSRP: Vl10 Grp 110 Standby router is local
*Jul 29 16:54:30.727: HSRP: Vl10 Grp 110 Speak -> Standby
*Jul 29 16:54:30.727: %HSRP-5-STATECHANGE: Vlan10 Grp 110 state Speak -> Standby
*Jul 29 16:54:30.728: HSRP: Peer not present
*Jul 29 16:54:30.728: HSRP: Vl10 Grp 110 Redundancy "hsrp-Vl10-110" state Speak -> Standby
*Jul 29 16:54:30.728: HSRP: Vl10 Grp 110 Hello  out 192.168.10.2 Standby pri 109 vIP 192.168.10.100
*Jul 29 16:54:31.082: HSRP: Vl10 Grp 110 Hello  in  192.168.10.1 Active  pri 110 vIP 192.168.10.100
*Jul 29 16:54:33.459: HSRP: Vl10 Grp 110 Hello  out 192.168.10.2 Standby pri 109 vIP 192.168.10.100
*Jul 29 16:54:33.811: HSRP: Vl10 Grp 110 Hello  in  192.168.10.1 Active  pri 110 vIP 192.168.10.100
*Jul 29 16:54:36.344: HSRP: Vl10 Grp 110 Hello  out 192.168.10.2 Standby pri 109 vIP 192.168.10.100
*Jul 29 16:54:36.378: HSRP: Vl10 Grp 110 Hello  in  192.168.10.1 Active  pri 110 vIP 192.168.10.100
*Jul 29 16:54:38.856: HSRP: Vl10 Grp 110 Hello  in  192.168.10.1 Active  pri 110 vIP 192.168.10.100
*Jul 29 16:54:38.876: HSRP: Vl10 Grp 110 Hello  out 192.168.10.2 Standby pri 109 vIP 192.168.10.100
*Jul 29 16:54:41.688: HSRP: Vl10 Grp 110 Hello  out 192.168.10.2 Standby pri 109 vIP 192.168.10.100
*Jul 29 16:54:41.717: HSRP: Vl10 Grp 110 Hello  in  192.168.10.1 Active  pri 110 vIP 192.168.10.100
```

E. Spanning Tree Troubleshooting

STP loop conditions or instability in a network can prevent proper communication of HSRP peers. Because of this improper communication, each peer becomes an active router. STP loops can cause broadcast storms, duplicated frames, and MAC table inconsistency. All of these problems affect the entire network, and especially HSRP. HSRP error messages can be the first indication of an STP issue.

When you troubleshoot STP, you *must* understand the STP topology of the network on each VLAN. You must determine what switch is the root bridge and which ports on the switch are on blocking and forwarding. Because each VLAN has its own STP topology, this information is very important on each VLAN.

1. Verify Spanning Tree Configuration

Be sure that STP is configured on every switch and bridging device in the network. Take note of where each switch believes the root bridge is located. Also, note the values of these timers:

- **Root Max Age**
- **Hello Time**
- **Forward Delay**

Issue the **show spanning-tree** command in order to see all of this information. By default, the command shows this information for all VLANs. But, you can also filter other VLAN information if you supply the VLAN number with the command. This information is very useful when you troubleshoot STP issues.

These three timers that you note in the **show spanning-tree** output are learned from the root bridge. These timers do not need to match the timers that are set on that specific bridge. But, be sure that the timers match

the root bridge in the case that this switch becomes the root bridge at any point. This match of the timers to the root bridge helps maintain continuity and ease of administration. The match also prevents a switch with incorrect timers from crippling the network.



Note: Enable STP for all VLANs at all times, regardless of whether there are redundant links in the network. If you enable STP in nonredundant networks, you prevent a breakage. A breakage can occur if someone bridges switches together with hubs or other switches and accidentally creates a physical loop. STP is also very useful in the isolation of specific problems. If the enablement of STP affects the operation of something in the network, there can be an existing problem that you need to isolate.

Here is sample output of the **show spanning-tree** command:

<#root>

L2Switch_1#

show spanning-tree vlan 10

```
VLAN0010
  Spanning tree enabled protocol rstp
  Root ID    Priority    32778
             Address     00fe.c8d3.8680
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec

  Bridge ID  Priority    32778  (priority 32768 sys-id-ext 10)
             Address     00fe.c8d3.8680
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  300 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
Gi1/0/3             Desg FWD 4         128.3    P2p
Gi1/0/10            Desg FWD 4         128.10   P2p Edge
Gi1/0/11            Desg FWD 4         128.11   P2p
Gi1/0/13            Desg FWD 4         128.13   P2p
Gi1/0/14            Desg FWD 4         128.14   P2p
Gi1/0/15            Desg FWD 4         128.15   P2p
Gi1/0/16            Desg FWD 4         128.16   P2p
Gi1/0/35            Desg FWD 4         128.35   P2p
```

L2Switch_1#

show spanning-tree vlan 11

```
VLAN0011
  Spanning tree enabled protocol rstp
  Root ID    Priority    32779
             Address     00fe.c8d3.8680
             This bridge is the root
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
```

```
  Bridge ID  Priority    32779  (priority 32768 sys-id-ext 11)
             Address     00fe.c8d3.8680
             Hello Time   2 sec  Max Age 20 sec  Forward Delay 15 sec
             Aging Time  300 sec

Interface           Role Sts Cost      Prio.Nbr Type
------------------- ---- --- --------- -------- --------------------------------
Gi1/0/3             Desg FWD 4         128.3    P2p
Gi1/0/10            Desg FWD 4         128.10   P2p Edge
Gi1/0/11            Desg FWD 4         128.11   P2p
Gi1/0/13            Desg FWD 4         128.13   P2p
Gi1/0/14            Desg FWD 4         128.14   P2p
Gi1/0/15            Desg FWD 4         128.15   P2p
Gi1/0/16            Desg FWD 4         128.16   P2p
Gi1/0/35            Desg FWD 4         128.35   P2p
```

Switch L2Switch_1 is the root of VLAN 10 and VLAN 11.

2. Spanning Tree Loop Conditions

In order for an STP loop to occur, there must be L2 physical redundancy in the network. An STP loop does not occur if there is no possibility of a physical loop condition. Symptoms of an STP loop condition are:

- Total network outage
- Loss of connectivity
- The report by network equipment of high process and system utilization

A single VLAN that experiences an STP loop condition can congest a link and starve the other VLANs of bandwidth. The **show interfaces <interface> controller** command notes which ports transmit or receive an excessive number of packets. Excessive broadcast and multicast can indicate ports that are part of an STP loop. Suspect an STP loop when broadcast or multicast traffic increases sharply beyond the established baseline and coincides with duplicate frames, MAC instability, topology changes, congestion, or high CPU utilization.



Note: The switch also counts STP bridge protocol data units (BPDUs) that are received and transmitted as multicast frames. A port that is in the STP blocking state still transmits and receives STP BPDUs.

<#root>

Router_2#

show interfaces gi1/0/1 controller

```
GigabitEthernet1/0/1 is up, line protocol is up (connected)
  Hardware is Gigabit Ethernet, address is 1880.90d8.5901 (bia 1880.90d8.5901)
  Description: PNP STARTUP VLAN
```

```
  MTU 1500 bytes, BW 1000000 Kbit/sec, DLY 10 usec,
     reliability 255/255, txload 1/255, rxload 1/255
  Encapsulation ARPA, loopback not set
  Keepalive set (10 sec)
  Full-duplex, 1000Mb/s, media type is 10/100/1000BaseTX
  input flow-control is on, output flow-control is unsupported
  ARP type: ARPA, ARP Timeout 04:00:00
  Last input 00:00:00, output 00:00:04, output hang never
  Last clearing of "show interface" counters never
  Input queue: 0/2000/0/0 (size/max/drops/flushes); Total output drops: 0
  Queueing strategy: fifo
  Output queue: 0/40 (size/max)
  5 minute input rate 33000 bits/sec, 31 packets/sec
  5 minute output rate 116000 bits/sec, 33 packets/sec
     9641686 packets input, 1477317083 bytes, 0 no buffer
     Received 1913802 broadcasts (1151766 multicasts)
     0 runts, 0 giants, 0 throttles
     0 input errors, 0 CRC, 0 frame, 0 overrun, 0 ignored
     0 watchdog, 1151766 multicast, 0 pause input
     0 input packets with dribble condition detected
     10702696 packets output, 4241534645 bytes, 0 underruns
     Output 3432 broadcasts (0 multicasts)
     0 output errors, 0 collisions, 2 interface resets
     9582 unknown protocol drops
     0 babbles, 0 late collision, 0 deferred
     0 lost carrier, 0 no carrier, 0 pause output
     0 output buffer failures, 0 output buffers swapped out
Transmit                  GigabitEthernet1/0/1          Receive
   4241534645 Total bytes                  1477317083 Total bytes
     10562003 Unicast frames                  7727884 Unicast frames
   4229489212 Unicast bytes                1291270617 Unicast bytes
      137261 Multicast frames                 1151766 Multicast frames
     11812065 Multicast bytes                91096867 Multicast bytes
         3432 Broadcast frames                  762036 Broadcast frames
      233368 Broadcast bytes                 94949599 Broadcast bytes
           0 System FCS error frames                0 IpgViolation frames
           0 MacUnderrun frames                     0 MacOverrun frames
           0 Pause frames                           0 Pause frames
           0 Cos 0 Pause frames                     0 Cos 0 Pause frames
           0 Cos 1 Pause frames                     0 Cos 1 Pause frames
           0 Cos 2 Pause frames                     0 Cos 2 Pause frames
           0 Cos 3 Pause frames                     0 Cos 3 Pause frames
           0 Cos 4 Pause frames                     0 Cos 4 Pause frames
           0 Cos 5 Pause frames                     0 Cos 5 Pause frames
           0 Cos 6 Pause frames                     0 Cos 6 Pause frames
           0 Cos 7 Pause frames                     0 Cos 7 Pause frames
           0 Oam frames                             0 OamProcessed frames
           0 Oam frames                             0 OamDropped frames
       38144 Minimum size frames              4165201 Minimum size frames
     4910833 65 to 127 byte frames            3126489 65 to 127 byte frames
     1237675 128 to 255 byte frames            750243 128 to 255 byte frames
     1029126 256 to 511 byte frames           1279281 256 to 511 byte frames
     2205966 512 to 1023 byte frames           103668 512 to 1023 byte frames
     1280952 1024 to 1518 byte frames          205229 1024 to 1518 byte frames
           0 1519 to 2047 byte frames           11575 1519 to 2047 byte frames
           0 2048 to 4095 byte frames               0 2048 to 4095 byte frames
           0 4096 to 8191 byte frames               0 4096 to 8191 byte frames
           0 8192 to 16383 byte frames              0 8192 to 16383 byte frames
           0 16384 to 32767 byte frame              0 16384 to 32767 byte frame
           0 > 32768 byte frames                    0 > 32768 byte frames
           0 Late collision frames                  0 SymbolErr frames
           0 Excess Defer frames                    0 Collision fragments
```

```
        0 Good (1 coll) frames                  0 ValidUnderSize frames
        0 Good (>1 coll) frames                 0 InvalidOverSize frames
        0 Deferred frames                       0 ValidOverSize frames
        0 Gold frames dropped                   0 FcsErr frames
        0 Gold frames truncated
        0 Gold frames successful
        0 1 collision frames
        0 2 collision frames
        0 3 collision frames
        0 4 collision frames
        0 5 collision frames
        0 6 collision frames
        0 7 collision frames
        0 8 collision frames
        0 9 collision frames
        0 10 collision frames
        0 11 collision frames
        0 12 collision frames
        0 13 collision frames
        0 14 collision frames
        0 15 collision frames
        0 Excess collision frames

LAST UPDATE 2384 msecs AGO
```

3. Topology Change Notification

Another command that is vital to the diagnosis of STP issues is the **show spanning-tree detail** command. This command tracks Topology Change Notification (TCN) messages back to the originator. These messages, sent as special BPDUs between switches, indicate that there has been a topology change on a switch. That switch sends a TCN out its root port. The TCN moves upstream to the root bridge. A bridge that receives a TCN BPDU acknowledges it by setting the Topology Change Acknowledgement (TCA) flag in a configuration BPDU on the port that received the TCN. If the receiving bridge is not the root bridge, it also sends a TCN BPDU through its root port. When the root bridge receives the TCN BPDU, it sets the Topology Change (TC) flag in the configuration BPDUs that it sends through the spanning-tree domain. This causes all nonroot bridges to set their MAC address table aging timer to the configuration STP forward delay.

In order to isolate this problem, access the root bridge for each VLAN and issue the **show spanning-tree <interface> detail** command for the switch-connected ports. The `last change occurred` entry gives the time that the last TCN was received. In this situation, you are too late to see who issued the TCNs that can have caused the possible STP loop. The **`Number of topology changes`** entry gives you an idea about the number of TCNs that occur. During an STP loop, this counter can increment every minute. Refer to Troubleshoot STP Problems & Related Design Considerations for more information.

Other useful information includes:

- Port of the last TCN
- Time of last TCN
- Current count of TCNs

Here is sample command output:

<#root>

L2Switch_1#

show spanning-tree vlan 10 detail

```
 VLAN0010 is executing the rstp compatible Spanning Tree protocol
  Bridge Identifier has priority 32768, sysid 10, address 00fe.c8d3.8680
  Configured hello time 2, max age 20, forward delay 15, transmit hold-count 6
  We are the root of the spanning tree
  Topology change flag not set, detected flag not set
  Number of topology changes 8 last change occurred 03:21:48 ago
          from GigabitEthernet1/0/35
  Times:  hold 1, topology change 35, notification 2
          hello 2, max age 20, forward delay 15
  Timers: hello 0, topology change 0, notification 0, aging 300

 Port 3 (GigabitEthernet1/0/3) of VLAN0010 is designated forwarding
   Port path cost 4, Port priority 128, Port Identifier 128.3.
   Designated root has priority 32778, address 00fe.c8d3.8680
   Designated bridge has priority 32778, address 00fe.c8d3.8680
   Designated port id is 128.3, designated path cost 0
   Timers: message age 0, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
   Link type is point-to-point by default
   BPDU: sent 6066, received 0

 Port 10 (GigabitEthernet1/0/10) of VLAN0010 is designated forwarding
   Port path cost 4, Port priority 128, Port Identifier 128.10.
   Designated root has priority 32778, address 00fe.c8d3.8680
   Designated bridge has priority 32778, address 00fe.c8d3.8680
   Designated port id is 128.10, designated path cost 0
   Timers: message age 0, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
   The port is in the portfast mode by portfast trunk configuration
   Link type is point-to-point by default
   BPDU: sent 6063, received 0

 Port 11 (GigabitEthernet1/0/11) of VLAN0010 is designated forwarding
   Port path cost 4, Port priority 128, Port Identifier 128.11.
   Designated root has priority 32778, address 00fe.c8d3.8680
   Designated bridge has priority 32778, address 00fe.c8d3.8680
   Designated port id is 128.11, designated path cost 0
   Timers: message age 0, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
   Link type is point-to-point by default
   BPDU: sent 6066, received 0

 Port 13 (GigabitEthernet1/0/13) of VLAN0010 is designated forwarding
   Port path cost 4, Port priority 128, Port Identifier 128.13.
   Designated root has priority 32778, address 00fe.c8d3.8680
   Designated bridge has priority 32778, address 00fe.c8d3.8680
   Designated port id is 128.13, designated path cost 0
   Timers: message age 0, forward delay 0, hold 0
   Number of transitions to forwarding state: 1
   Link type is point-to-point by default
```

```
  BPDU: sent 6066, received 3

Port 14 (GigabitEthernet1/0/14) of VLAN0010 is designated forwarding
  Port path cost 4, Port priority 128, Port Identifier 128.14.
  Designated root has priority 32778, address 00fe.c8d3.8680
  Designated bridge has priority 32778, address 00fe.c8d3.8680
  Designated port id is 128.14, designated path cost 0
  Timers: message age 0, forward delay 0, hold 0
  Number of transitions to forwarding state: 1
  Link type is point-to-point by default
  BPDU: sent 6066, received 3

Port 15 (GigabitEthernet1/0/15) of VLAN0010 is designated forwarding
  Port path cost 4, Port priority 128, Port Identifier 128.15.
  Designated root has priority 32778, address 00fe.c8d3.8680
  Designated bridge has priority 32778, address 00fe.c8d3.8680
  Designated port id is 128.15, designated path cost 0
  Timers: message age 0, forward delay 0, hold 0
  Number of transitions to forwarding state: 1
  Link type is point-to-point by default
  BPDU: sent 6067, received 0

Port 16 (GigabitEthernet1/0/16) of VLAN0010 is designated forwarding
  Port path cost 4, Port priority 128, Port Identifier 128.16.
  Designated root has priority 32778, address 00fe.c8d3.8680
  Designated bridge has priority 32778, address 00fe.c8d3.8680
  Designated port id is 128.16, designated path cost 0
  Timers: message age 0, forward delay 0, hold 0
  Number of transitions to forwarding state: 1
  Link type is point-to-point by default
  BPDU: sent 6067, received 0

Port 35 (GigabitEthernet1/0/35) of VLAN0010 is designated forwarding
  Port path cost 4, Port priority 128, Port Identifier 128.35.
  Designated root has priority 32778, address 00fe.c8d3.8680
  Designated bridge has priority 32778, address 00fe.c8d3.8680
  Designated port id is 128.35, designated path cost 0
  Timers: message age 0, forward delay 0, hold 0
  Number of transitions to forwarding state: 1
  Link type is point-to-point by default
  BPDU: sent 6067, received 0
```

This output shows that the last topology change occurred from device connected off interface GigabitEthernet1/0/35. Next, issue the same **show spanning-tree detail** command from this device in order to try to track the issue. If this switch that generates the TCNs is only attached to PC or endpoints, be sure that STP PortFast is enabled on these ports. STP PortFast suppresses STP TCNs when a port transitions states.

Refer to these documents for information about STP and how to troubleshoot link transitions that are associated with network interface cards (NICs):

- Using PortFast and Other Commands to Fix Workstation Startup Connectivity Delays
- Understand Rapid Spanning Tree Protocol (802.1w)
- Troubleshoot STP Problems & Related Design Considerations

4. Broadcast Suppression

Enable broadcast suppression in order to help cut down the impact from a broadcast storm. A broadcast storm is one of the main side effects of an STP loop. Here is sample command output.



Note: Consider storm control in order to limit the impact of a broadcast, multicast, or unknown-unicast storm. Set thresholds from measured normal traffic and validate the configured action before deployment.

<#root>

L2Switch_1#

show run interface TenGigabitEthernet1/1/5

```
Building configuration...

Current configuration : 279 bytes
!
interface TenGigabitEthernet1/1/5
 switchport trunk allowed vlan 300-309
 switchport mode trunk
 storm-control broadcast level 30.00
 storm-control multicast level 30.00
 storm-control unicast level 30.00
 spanning-tree guard root
end

L2Switch_1#
```

show storm-control broadcast

```
Key: U - Unicast, B - Broadcast, M - Multicast
Interface  Filter State   Upper        Lower        Current     Action     Type
---------  -------------  -----------  -----------  ----------  ---------  ----
Te1/1/5    Forwarding          30.00%       30.00%       0.00%  None       B
Te1/1/7    Link Down           30.00%       30.00%       0.00%  None       B
Te1/1/8    Forwarding          10.00%       10.00%       0.00%  None       B

L2Switch_1#show storm-control multicast
Key: U - Unicast, B - Broadcast, M - Multicast
Interface  Filter State   Upper        Lower        Current     Action     Type
---------  -------------  -----------  -----------  ----------  ---------  ----
Te1/1/5    Forwarding          30.00%       30.00%       0.00%  None       M
Te1/1/7    Link Down           30.00%       30.00%       0.00%  None       M
```

5. Console and Telnet Access

Console or Telnet traffic to the switch often becomes too sluggish to properly track down an offending device during an STP loop. During an approved maintenance window and with out-of-band access available, isolate suspected redundant links one at a time after verifying that the remaining path preserves required connectivity. Permit STP to reconverge and verify stability before restoring each link. If the STP loop returns after you add one particular segment, you have identified the offending devices.

6. Spanning Tree Features: Portfast, UplinkFast, and BackboneFast

Verify that PortFast, UplinkFast, and BackboneFast are configured properly. When you troubleshoot STP issues, verify whether legacy UplinkFast or BackboneFast is configured and whether the active spanning-tree mode uses those features. Change the configuration only during a controlled maintenance window after an impact assessment. In addition, verify that STP PortFast is only enabled on ports that are directly connected to nonbridging hosts. Nonbridging hosts include user workstations and routers without bridge groups. Do not enable PortFast on ports that are connected to hubs or other switches. Here are some documents to help understand and configure these features:

Understand the Spanning Tree PortFast BPDU Guard Enhancement

Understand and Configure the Cisco UplinkFast Feature

7. BPDU Guard

When you enable PortFast BPDU guard, a nontrunking, PortFast-enabled port is moved into an errdisable state at the receipt of a BPDU on that port. This feature helps you find ports that are incorrectly configured for PortFast. The feature also detects where devices reflect packets or inject STP BPDUs into the network. When you troubleshoot STP issues, you can enable this feature to help isolate the STP issue.

<#root>

L2Switch_1#

configure terminal

Enter configuration commands, one per line. End with CNTL/Z.
L2Switch_1(config)#

spanning-tree portfast bpduguard

L2Switch_1(config)#

end

8. VTP Pruning

When VTP Pruning is enabled in the network, it can cause the devices of an HSRP group to go active. This results in IP conflicts among the gateways and cause traffic issues. Ensure that VTP does not prune the VLAN of any HSRP group from required trunks in the network.

F. Divide and Conquer

If all other attempts to isolate or resolve HSRP fail, the "divide and conquer" method is the next approach. This method helps isolate the network and components that make up the network. Divide and conquer involves any one of the guidelines in this list:

Note: This list repeats some guidelines from other sections of this document.

- Create a test VLAN for HSRP and isolated VLAN to switch with HSRP routers.
- Disconnect all redundant ports.
- Break FEC ports into single connected ports.
- Reduce HSRP group members to only two members.
- Prune trunk ports such that only necessary VLANs propagate across those ports.
- Disconnect connected switches in the network until the problems cease.

Related Information

- LAN Switching
- Technical Support & Documentation - Cisco Systems