



# Chapter 4: Link aggregation & First-hop redundancy protocols

CCNA Routing and Switching

Scaling Networks

*With additional content from Overview of the IEEE 802.3ad-2000 (clause 43) by Anthony Mwingira, Jon Sterritt of Bridge Functions Consortium*



# Chapter 4 - Sections & Objectives

## ■ 4.1 Link Aggregation Concepts

- Explain link aggregation operation in a switched LAN environment.
  - Describe link aggregation.
  - Describe EtherChannel technology.

## ■ 4.2 Link Aggregation Configuration

- Implement link aggregation to improve performance on high-traffic switch links.
  - Configure link aggregation.
  - Troubleshoot a link aggregation implementation.


## ■ 4.3 First Hop Redundancy Protocols

- Implement HSRP
  - Explain the purpose and operation of first hop redundancy protocols.
  - Explain how HSRP operates.
  - Configure HSRP using Cisco IOS commands.
  - Troubleshoot HSRP.

# 4.1 Link Aggregation Concepts

# Introduction to Link Aggregation


- It is possible to combine the number of physical links between switches to increase the overall speed of switch-to-switch communication.
- STP will block redundant links to prevent routing loops.



# Introduction to Link Aggregation

- Link Aggregation (LA):
  - The combination of multiple physical links to function as a single logical link.
- Link Aggregation Group (LAG):
  - A group of physical links that make up a single aggregation.
  - Each physical port may be a member of a single LAG.
- Link Aggregation Control Protocol (LACP):
  - The protocol spec



tion.



LAG formed using three links directly connected between two adjacent devices

# Goals of Link Aggregation

- Increased bandwidth
  - multiple links combined into one logical link
- Linearly incremental bandwidth
  - increase in unit multiples
- Increased availability
  - failure of a single link within aggregation need not cause MAC client failure
- Load sharing
  - MAC Client traffic may be distributed across links
- Auto configuration
  - in even of changes in physical connection
- Rapid configuration and reconfiguration
  - also in the event of changes in physical connection



# Goals of Link Aggregation

- Deterministic behavior
  - independent of algorithm chosen (implementer can choose internal algorithm)
  - the configuration can be made to resolve deterministically, i.e. resulting aggregation can be made independent of the order in which events occur, and be completely determined by the capabilities of the individual links and their physical connectivity.
- Low risk of duplication or mis-ordering
  - during link reconfiguration or steady-state operation there is a high probability that frames are neither duplicated or mis-ordered.
- Backwards compatibility with aggregation-unaware devices
  - link that can not take part in an aggregation operate as normal, individual links.
- Multipoint aggregations
  - does not support aggregations among more than two systems;
  - multiple aggregations are allowed, but 3 systems cannot share a single Link Aggregation Group.

# Principles of Link Aggregation

- Link Aggregation allows a MAC Client to treat a set of one or more ports as if it were a single port.
- A Mac Client communicates with a set of ports through an Aggregator. An Aggregator binds to one or more ports within a System.
- It is the responsibility of an Aggregator to distribute frame transmissions from the MAC Client to the various ports, and to collect received frames from the ports and pass them to the MAC Client transparently.
- A given port will bind to (at most) a single Aggregator at any time. A MAC Client is served by a single Aggregator at a time.
- The binding of ports to Aggregators within a System is managed by the Link Aggregation Control function for that System.



# Principles of Link Aggregation

- Binding may be under manual control through direct manipulation of the state variables of Link Aggregation (e.g. Keys) by net manager. Automatic determination, configuration, binding, and monitoring may occur through the use of Link Aggregation Control Protocol (LACP)
- Frame ordering must be maintained. The distributor ensures that all frames of a given conversation are past to a single port. They are to be passed in order they were received from that port.
- Conversations may be moved among ports within aggregation, both for load balancing and to maintain availability in the event of a link failures.
- Each port is assigned a unique, globally administered MAC address. (Prospective of a MAC Client)
- Each Aggregator is assigned a unique, globally administered MAC address.

# Link Aggregation Group (LAG) Operation

- Frame Distribution
  - Responsible for taking frames from MAC Client and submitting them for transmission on appropriate ports.
- Frame Collection
  - Responsible for passing frames received from the various ports to the MAC Client
- Aggregator Parser / Multiplexers
  - When transmitting, pass frame requests to the appropriate port.
  - When receiving, distinguish frames and pass them to appropriate entity.
- Aggregator
  - This is the combination of frame Distribution and Collection, and Aggregator Parser / Multiplexers.

# Link Aggregation Group (LAG) Operation

- Aggregation Control
  - Responsible for configuration and control of the Link Aggregation.
- Control Parser/Multiplexers
  - On transmit, pass frame transmission requests from the Aggregator and Control entities to the appropriate port.
  - On receive, distinguish Link Aggregation Control PDUs from other frames, and passing LACPDU to the appropriate sub-layer entity and all other frames to the Aggregator.

# Configuring Link Aggregation Group (LAG)

- Different Strategies by Vendor
- Cisco
  - go into the interface prompt and put that interface into a LAG, i.e., "channel-group"
- Dell
  - command line, almost identical to Cisco, Web Interface the ports are configured from the LAG menu
- HP, Nortel
  - LAG's configured as "trunks" that run LACP as a protocol.
- 3Com
  - From menu driven command line ports must be activated for LACP (port menu) and then added to an activated LAG (Link Agg. Menu).
- Other "odd" implementations
  - require setting keys, most devices require ports assigned to a LAG, but some require setting the Key for each port, the result is equivalent.

# Link Aggregation Control Protocol (LACP)

- Link Aggregation Control configures and controls the Link Aggregation sub layer using static information local to the function and dynamic information exchanged by means of the **Link Aggregation Control Protocol**.
- The 802.3ad IEEE standard presents the means for the forming of a single Ethernet link automatically from two or more Ethernet links using LACP.
- In addition to the IEEE standards, there are also proprietary types of LACP developed by companies like Cisco.

# Link Aggregation Control Protocol (LACP)


- For each aggregatable port in the system:
  - Maintains configuration information to control aggregation
  - Exchange configuration information with other with other systems to allocate the link to a Link Aggregation Group. *A given link is allocated to, at most, one Link Aggregation Group at a time.*
  - Attaches the port to the Aggregator used by the Link Aggregator Group, and detaches the port from the Aggregator when it is no longer used by the Group.
  - Uses information from the Partner Systems' Link Aggregation Control entity to enable or disable the Aggregator or Distributor.
  - Checking that candidate links are actually be aggregated.
  - Controlling link addition to Link Aggregator Group and the creation of a Group if necessary
  - Monitoring the status of aggregated links.
  - Removing a link from a Link Aggregator Group if its membership is no longer valid, and removing the group if it no longer has any member links.

# Proprietary Link Aggregation Standards

- Cisco
  - EtherChannel and Port Aggregation Protocol
- Juniper
  - Aggregated Ethernet
- ZTE
  - Smartgroup
- Huawei
  - Eth-Trunk
- AVAYA
  - Multi-Link Trunking, Split Multi-Link Trunking, Routed Split Multi-Link Trunking and Distributed Split Multi-Link Trunking

# EtherChannel

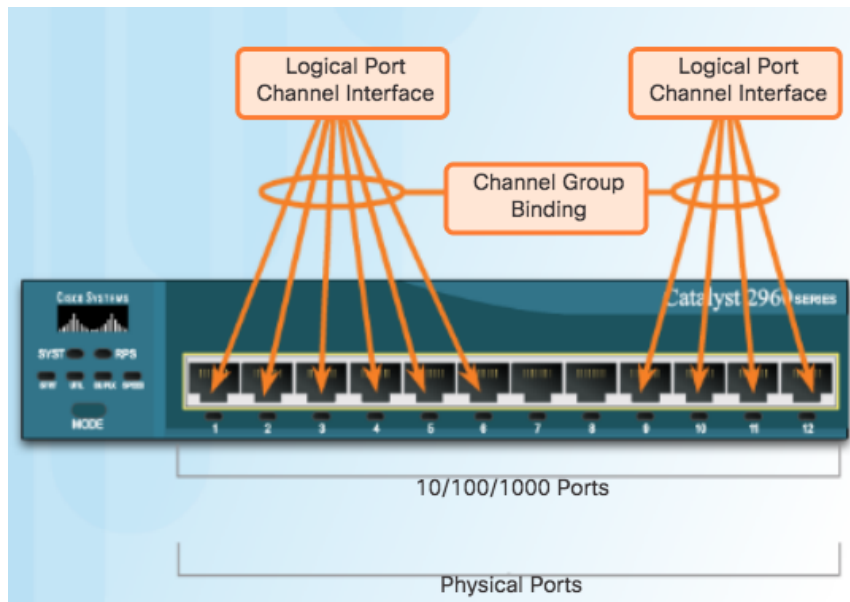
- EtherChannel is a Cisco specific link aggregation standard that groups multiple physical ports into one or more logical EtherChannel links.
- EtherChannels can be formed by using Cisco PAgP or LACP protocol.
- Advantages of EtherChannel:
  - Most configuration tasks can be done on the EtherChannel interface instead of on each individual port.
  - EtherChannel relies on existing switch ports.
  - Load balancing takes place between links that are part of the same EtherChannel.
  - EtherChannel creates an aggregation that is seen as one logical link.
  - EtherChannel provides redundancy because the overall link is seen as one logical connection.





# EtherChannel

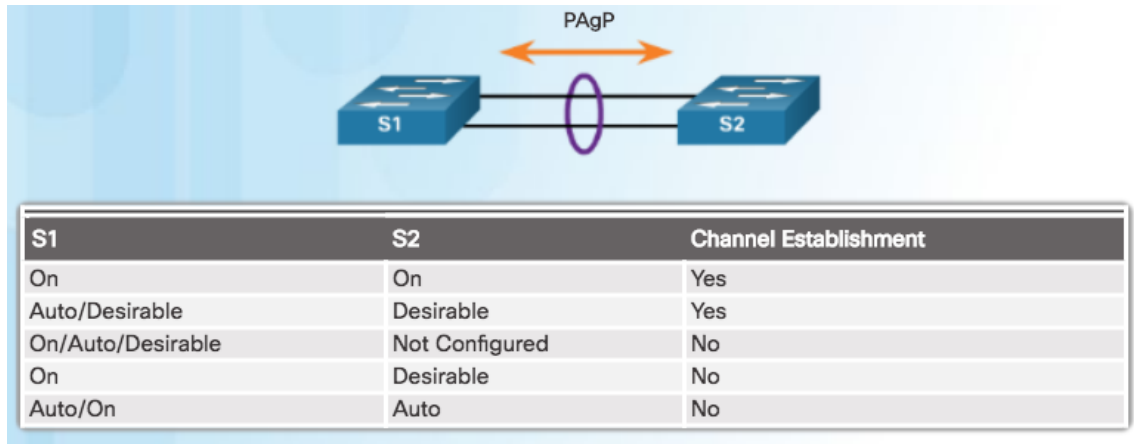
- EtherChannel Restrictions:
  - **Interface types cannot be mixed.** (Fast Ethernet + Gigabit Ethernet cannot be grouped.)
  - Provides full-duplex bandwidth up to 800 Mbps (Fast EtherChannel) or 8 Gbps (Gigabit EtherChannel)
  - Cisco IOS Switch can support 6 EtherChannels.
  - Created between two switches or a server and switch.
  - If one side is configured as trunk, the other side must be a trunk within same native VLAN.
  - Each EtherChannel has a logical port channel interface and changes to a channel affects its physical interfaces.



# EtherChannel Operation

## Port Aggregation Protocol

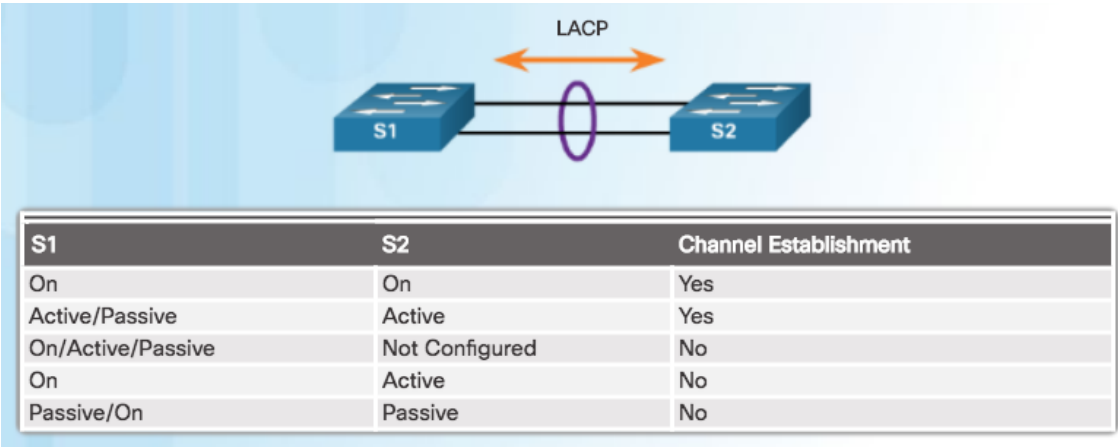
- PAgP (“Pag-P”) Cisco-proprietary protocol
- PAgP modes:
  - On: Channel member without negotiation (no protocol).
  - Desirable: Actively asking if the other side can or will participate.
  - Auto: Passively waiting for the other side.



# EtherChannel Operation

## Link Aggregation Control Protocol

- LACP multivendor environment
- LACP modes:
  - On: Channel member without negotiation (no protocol).
  - Active: Actively asking if the other side can or will participate.
  - Passive: Passively waiting for the other side.




## 4.2 Link Aggregation Configuration

# Configuring EtherChannel

## Configuration Guidelines

- Configuration Settings Match on Both Switches
  - Same speed and duplex mode.
  - All interfaces in a bundle must be assigned to the same VLAN, or configured as a trunk.
  - Trunk must support same range of VLANs.
- If Configuration Settings Do Not Match
  - EtherChannel not formed between S1 and S2
  - **Note: When changing settings, configure them in port channel interface configuration mode. The configuration applied to the port channel interface also affects the individual interfaces.**




S1 Port Configurations

Speed	1 Gb/s
Duplex	Full
VLAN	10

S2 Port Configurations

Speed	1 Gb/s
Duplex	Full
VLAN	10



S2 Port Configurations

Speed	1 Gb/s
Duplex	Half
VLAN	10

S2 Port Configurations


Speed	1 Gb/s
Duplex	Full
VLAN	10

# Configuring EtherChannel

## Configuring Interfaces

- This configuration creates EtherChannel with LACP and configures trunking.
  - Step 1: Specify the interfaces that compose the EtherChannel group.
  - Step 2: Create the port channel interface with the **channel-group** command in **active** mode. (Channel group number needs to be selected.)
  - Step 3: Change Layer 2 settings in port channel interface configuration mode.

```
S1(config)# interface range FastEthernet0/1 - 2
S1(config-if-range)# channel-group 1 mode active
Creating a port-channel interface Port-channel 1
S1(config-if-range)# interface port-channel 1
S1(config-if)# switchport mode trunk
S1(config-if)# switchport trunk allowed vlan 1,2,20
```



# Verifying EtherChannel

Verifies the interface status.

```
S1# show interfaces port-channel 1
Port-channel1 is up, line protocol is up (connected)
  Hardware is EtherChannel, address is 0cd9.96e8.8a02 (bia 0cd9.96e8.8a02)
    MTU 1500 bytes, BW 200000 Kbit/sec, DLY 100 usec,
      reliability 255/255, txload 1/255, rxload 1/255
    Encapsulation ARPA, loopback not set
<Output omitted>
```

Displays a one-line summary per channel group. SU indicates in use.

```
S1# 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
```

```
Number of channel-groups in use: 1
Number of aggregators:           1
```

Group	Port-channel	Protocol	Ports
1	Po1(SU)	LACP	Fa0/1(P) Fa0/2(P)

Displays port channel information.

```
S1# show etherchannel port-channel
Channel-group listing:
-----
```

```
Group: 1
-----
Port-channels in the group:
-----
```

```
Port-channel: Po1      (Primary Aggregator)
-----
```

```
Age of the Port-channel   = 0d:06h:23m:49s
Logical slot/port         = 2/1           Number of ports = 2
HotStandBy port = null
Port state                 = Port-channel Ag-Inuse
Protocol                   = LACP
Port security              = Disabled
```

## Ports in the Port-channel:

Index	Load	Port	EC state	No of bits
0	55	Fa0/1	Active	4
1	45	Fa0/2	Active	4

```
Time since last port bundled: 0d:05h:52m:59s Fa0/2
Time since last port Un-bundled: 0d:05h:53m:05s Fa0/2
```

# Verifying EtherChannel

Displays role of particular interface in an EtherChannel.

**SI# show interfaces f0/1 etherchannel**

Port state = Up Mstr Assoc In-Bndl  
Channel group = 1 Mode = Active Gcchange = -  
Port-channel = Po1 GC = - Pseudo port-channel = Po1  
Port index = 0 Load = 0x00 Protocol = LACP

Flags: S - Device is sending Slow LACPDUs F - Device is sending fast LACPDUs.  
A - Device is in active mode. P - Device is in passive mode.

Local information:

Port	Flags	State	LACP port Priority	Admin Key	Oper Key	Port Number	Port State
Fa0/1	SA	bndl	32768	0x1	0x1	0x102	0x3D

Partner's information:

Port	Flags	LACP port Priority	Dev ID	Age	Admin key	Oper Key	Port Number	Port State
Fa0/1	SA	32768	0cd9.96d2.4000	13s	0x0	0x1	0x102	0x3D

Age of the port in the current state: 0d:06h:06m:51s



# Troubleshooting EtherChannel

- All interfaces within EtherChannel must have the same:
  - speed
  - duplex mode
  - native and allowed VLANs on trunk (Ports with different native VLANs cannot form an EtherChannel.)
  - assigned to same VLAN

Output indicates that the EtherChannel is down (SD).

```
S1# 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
```

```
Number of channel-groups in use: 1
Number of aggregators:           1
```

Group	Port-channel	Protocol	Ports	
1	Pol(SD)	-	Fa0/1(D)	Fa0/2(D)

# Troubleshooting EtherChannel

Incompatible PAgP modes configured on S1 and S2.

```
S1# show run | begin interface port-channel
interface Port-channel1
  switchport mode trunk
!
interface FastEthernet0/1
  switchport mode trunk
  channel-group 1 mode on
!
interface FastEthernet0/2
  switchport mode trunk
  channel-group 1 mode on
!
<Output omitted>
S2# show run | begin interface port-channel
interface Port-channel1
  switchport mode trunk
!
interface FastEthernet0/1
  switchport mode trunk
  channel-group 1 mode desirable
!
interface FastEthernet0/2
  switchport mode trunk
  channel-group 1 mode desirable
```

S1	S2	Channel Establishment
On	On	Yes
Auto/Desirable	Desirable	Yes
On/Auto/Desirable	Not Configured	No
On	Desirable	No
Auto/On	Auto	No

PAgP mode on the EtherChannel is changed to desirable and the EtherChannel becomes active.

```
S1(config)# no interface port-channel 1
S1(config)# interface range f0/1 - 2
S1(config-if-range)# channel-group 1 mode desirable
Creating a port-channel interface Port-channel 1

S1(config-if-range)# no shutdown
S1(config-if-range)# interface port-channel 1
S1(config-if)# switchport mode trunk
S1(config-if)# end
S1# 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
```

Number of channel-groups in use: 1			
Number of aggregators: 1			
Group Port-channel Protocol Ports			
-----+-----+-----+-----			
1	Pol(SU)	PAgP	Fa0/1(P) Fa0/2(P)

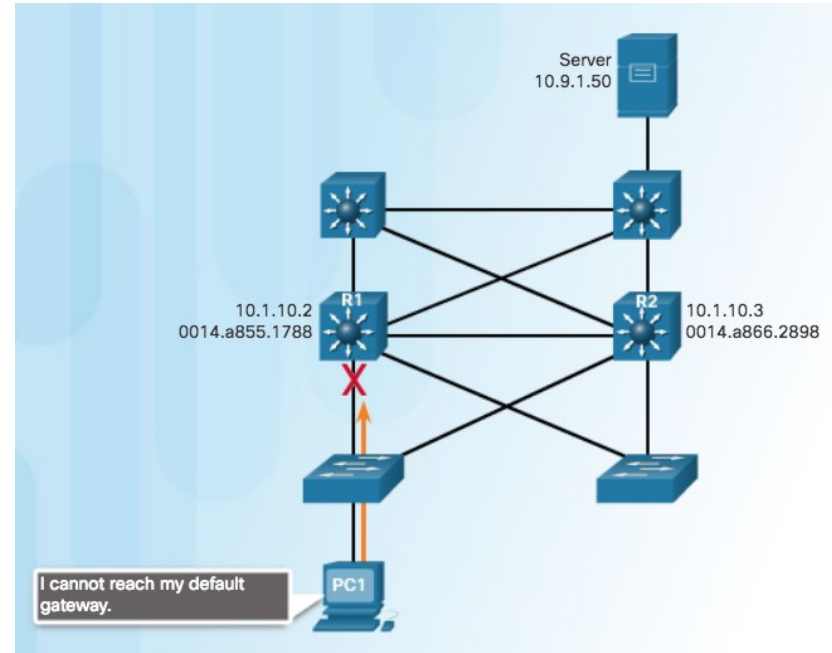
S1	S2	Channel Establishment
On	On	Yes
Auto/Desirable	Desirable	Yes
On/Auto/Desirable	Not Configured	No
On	Desirable	No
Auto/On	Auto	No

# 4.3 First Hop Redundancy Protocols

# Concept of First Hop Redundancy Protocols

## Default Gateway Limitations

- A mechanism is needed to provide alternate default gateways in switched networks where two or more routers are connected to the same VLANs.
  - Note: In the graphic, a multilayer switch is acting as the default gateway and used for routing.
  - In a switched network, each client receives only one default gateway.
  - There is no way to use a secondary gateway, even if a second path exists to carry packets off the local segment.
  - In the figure, R1 is responsible for routing packets from PC1. If R1 becomes unavailable, R2 can route packets that would have gone through R1.




- End devices are typically configured with a single IP address for a default gateway.
- If that default gateway IP address cannot be reached, the local device is unable to send packets off the local network.

# Concept of First Hop Redundancy Protocols

## Router Redundancy

- To prevent a single point of failure at the default gateway, implement a virtual router.
  - Present the illusion of a single router to the hosts on the LAN.
  - By sharing an IP address and a MAC address, two or more routers can act as a single virtual router.
  - IPv4 address of the virtual router is configured as the default gateway for the workstations on a specific IPv4 segment.
  - ARP resolution returns the MAC address of the virtual router.
  - Physical router that forwards traffic is transparent to the host devices.




- A redundancy protocol provides the mechanism for determining which router should take the active role in forwarding traffic.
- Ability of a network to dynamically recover from the failure of a device acting as a default gateway is known as **first-hop redundancy**.

# Concept of First Hop Redundancy Protocols

## Steps for Router Failover


- When the active router fails, the redundancy protocol transitions the standby router to the new active router role.
- These are the steps that take place when the active router fails:
  1. The standby router stops seeing hello messages from the forwarding router.
  2. The standby router assumes the role of the forwarding router.
  3. Because the new forwarding router assumes both the IPv4 and MAC addresses of the virtual router, the host devices see no disruption in service.



# Concept of First Hop Redundancy Protocols

## First Hop Redundancy Protocols

- **Hot Standby Router Protocol (HSRP)** - A Cisco-proprietary FHRP designed to allow for transparent failover of a first-hop IPv4 device.
  - Active device is the device that is used for routing packets.
  - Standby device is the device that takes over when the active device fails.
  - Function of the HSRP standby router is to monitor the operational status of the HSRP group and to quickly assume packet-forwarding responsibility if the active router fails.
- **HSRP for IPv6** - Cisco-proprietary FHRP providing the same functionality of HSRP, but in an IPv6 environment.



- HSRP defines a group of routers - one active and one standby.
- Virtual IP and MAC addresses are shared between the two routers.
- To verify HSRP state, use the **show standby** command.
- HSRP is Cisco proprietary.
- VRRP is a standard protocol.


# First Hop Redundancy Protocols (Cont.)

- **Gateway Load Balancing Protocol (GLBP)** - A more recent Cisco-proprietary FHRP that protects data traffic from a failed router or circuit allowing **load balancing** between a group of redundant routers.
- **GLBP for IPv6** - Cisco-proprietary FHRP providing the same functionality of GLBP.
- **Virtual Router Redundancy Protocol version 2** - A nonproprietary/open (albeit patent encumbered) protocol that dynamically assigns responsibility for one or more virtual routers to the VRRP routers on an IPv4 LAN.
  - One router is elected as the virtual router master, with the other routers acting as backups, in case the virtual router master fails.
- **VRRPv3** - Capability to support IPv4 and IPv6.
- **Common Address Redundancy Protocol (CARP)** - A patent-free unencumbered alternative to Cisco's HSRP developed in October 2003.



# HSRP Operations


## HSRP Overview



- One of the routers is selected by HSRP to be the active router and default gateway.
- Other router will become the standby router.
- If active router fails, standby assumes the role of active router and default gateway.
- Hosts are configured with single default gateway VIRTUAL address that is recognizable by both the active and standby routers.

# HSRP Operations


## HSRP Versions



Version	HSRP V1 (Default)	HSRP V2
Group numbers	0 to 255	0 to 4095
Multicast address	224.0.0.2	224.0.0.102 or FF02::66
Virtual MAC address	0000.0C07.AC00 - 0000.0C07.ACFF (last two digits group number)	<b>IPv4</b> 0000.0C9F.F000 to 0000.0C9F.FFFF <b>IPv6</b> 0005.73A0.0000-0005.73A0.0FFF (last three digits group number)
Support for MD5 authentication	No	Yes

# HSRP Priority and Preemption

- Role of active and standby routers determined by election process.
- By default, the router with the numerically highest IPv4 address is elected as the active router.
- Control HSRP election with priority and do not use highest address.
- HSRP Priority
  - Used to determine active router.
  - Default HSRP priority is 100.
  - Range is 0 to 255 and router with highest priority will become active.
  - Use the **standby priority interface** command.
- HSRP Preemption
  - Preemption - ability of HSRP router to trigger the re-election process.
  - To force a new HSRP election process, preemption must be enabled using **standby preempt interface**.
  - A router that comes online with the a higher priority will become the active router.



# HSRP States and Timers

- The active and standby HSRP routers send hello packets to the HSRP group multicast address every 3 seconds, by default. The standby router will become active if it does not receive a hello message from the active router after 10 seconds.
- You can lower these timer settings to speed up the failover or preemption. However, to avoid increased CPU usage and unnecessary standby state changes, do not set the hello timer below 1 second or the hold timer below 4 seconds.

State	Definition
Initial	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 a hello message from the active router. In this state, the router 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.
Standby	The router is a candidate to become the next active router and sends periodic hello messages.
Active	The router currently forwards packets that are sent to the group virtual MAC address. The router sends periodic hello messages.

# HSRP Configuration


## HSRP Configuration Commands

- Step 1. Configure HSRP version 2.
- Step 2. Configure the virtual IP address for the group.
- Step 3. Configure the priority for the desired active router to be greater than 100.
- Step 4. Configure the active router to preempt the standby router in cases where the active router comes online after the standby router.

Command	Definition
Router(config-if)# <b>standby version 2</b>	Configures HSRP to use version 2. HSRP version 1 is the default.
Router(config-if)# <b>standby</b> [group-number] ip-address	Configures the HSRP virtual IP address that will be used by the specified group. If no group is configured, then the virtual IP address is assigned to group 0.
Router(config-if)# <b>standby</b> [group-number] priority [priority-value]	Configures the desired active router with a higher priority than default priority of 100. Range is 0 to 255. If no priority is configured or if priority is equal, than the router with the highest IP address has priority.
Router(config-if)# <b>standby</b> [group-number] preempt	Configures a router to preempt the currently active router.

# HSRP Configuration

## HSRP Sample Configuration



```
R1(config)# interface g0/1
R1(config-if)# ip address 172.16.10.2 255.255.255.0
R1(config-if)# standby version 2
R1(config-if)# standby 1 ip 172.16.10.1
R1(config-if)# standby 1 priority 150
R1(config-if)# standby 1 preempt
R1(config-if)# no shutdown
!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!!
R2(config)# interface g0/1
R2(config-if)# ip address 172.16.10.3 255.255.255.0
R2(config-if)# standby version 2
R2(config-if)# standby 1 ip 172.16.10.1
R2(config-if)# no shutdown
```

# HSRP Configuration


## HSRP Verification

<pre>R1# show standby GigabitEthernet0/1 - Group 1 (version 2)   State is Active     5 state changes, last state change 01:02:18   Virtual IP address is 172.16.10.1   Active virtual MAC address is 0000.0c9f.f001     Local virtual MAC address is 0000.0c9f.f001 (v2 default)   Hello time 3 sec, hold time 10 sec     Next hello sent in 1.120 secs   Preemption enabled   Active router is local   Standby router is 172.16.10.3, priority 100 (expires in 9.392 sec)   Priority 150 (configured 150)   Group name is "hsrp-Gi0/1-1" (default)</pre>	<pre>R1# show standby brief P indicates configured to preempt.   Interface  Grp  Pri P State  Active      Standby      Virtual IP Gi0/1      1    150 P Active local    172.16.10.3  172.16.10.1</pre>
<pre>R2# show standby GigabitEthernet0/1 - Group 1 (version 2)   State is Standby     5 state changes, last state change 01:03:59   Virtual IP address is 172.16.10.1   Active virtual MAC address is 0000.0c9f.f001     Local virtual MAC address is 0000.0c9f.f001 (v2 default)   Hello time 3 sec, hold time 10 sec     Next hello sent in 0.944 secs   Preemption disabled   Active router is 172.16.10.2, priority 150 (expires in 8.160 sec)     MAC address is fc99.4775.c3e1   Standby router is local   Priority 100 (default 100)   Group name is "hsrp-Gi0/1-1" (default)</pre>	<pre>R2# show standby brief P indicates configured to preempt.   Interface  Grp  Pri P State  Active      Standby      Virtual IP Gi0/1      1    100 Standby 172.16.10.2 local    172.16.10.1</pre>

# HSRP Troubleshooting

## HSRP Failure

- Most issues will arise during one of the following HSRP functions:
  - Failing to successfully elect the active router that controls the virtual IP for the group
  - Failure of the standby router to successfully keep track of the active router
  - Failing to determine when control of the virtual IP for the group should be handed over to another router
  - Failure of end devices to successfully configure the virtual IP address as the default gateway





# Common HSRP Configuration Issues

Use the debug commands to detect common configuration issues:

- HSRP routers are not connected to the same network segment. Although this could be a physical layer issue, it could also be a VLAN subinterface configuration issue.
- HSRP routers are not configured with IPv4 addresses from the same subnet. HSRP hello packets are local. They are not routed beyond the network segment. Therefore, a standby router would not know when the active router fails.
- HSRP routers are not configured with the same virtual IPv4 address. The virtual IPv4 address is the default gateway for end devices.
- HSRP routers are not configured with the same HSRP group number. This will cause each router to assume the active role.
- End devices are not configured with the correct default gateway address. Although not directly related to HSRP, configuring the DHCP server with one of the real IP addresses of the HSRP router would mean that end devices would only have connectivity to remote networks when that HSRP router is active.

# HSRP Troubleshooting


## HSRP Debug Commands

```
R2# debug standby ?
errors    HSRP errors
events    HSRP events
packets   HSRP packets terse
Display limited range of HSRP errors, events and packets
<cr>
```

View the HSRP Hello Packets on Standby Router

```
R2# debug standby packets
*Dec  2 15:20:12.347: HSRP: Gi0/1 Grp 1 Hello in 172.16.10.2
Active pri 150 vIP 172.16.10.1

*Dec  2 15:20:12.643: HSRP: Gi0/1 Grp 1 Hello out 172.16.10.3
Standby pri 100 vIP 172.16.10.1
```




# HSRP Troubleshooting

## HSRP Debug Commands (Cont.)

R1 Fails and R2 is Elected Active HSRP Router

```
!!!!!!R1 is powered off!!!!
R2# debug standby terse
HSRP:
  HSRP Errors debugging is on
  HSRP Events debugging is on
    (protocol, neighbor, redundancy, track, arp, interface)
  HSRP Packets debugging is on
    (Coup, Resign)
R2#
*Dec  2 16:11:31.855: HSRP: Gi0/1 Grp 1 Standby: c/Active timer expired
  (172.16.10.2)
*Dec  2 16:11:31.855: HSRP: Gi0/1 Grp 1 Active router is local, was
  172.16.10.2
*Dec  2 16:11:31.855: HSRP: Gi0/1 Nbr 172.16.10.2 no longer active for
  group 1 (Standby)
*Dec  2 16:11:31.855: HSRP: Gi0/1 Nbr 172.16.10.2 Was active or standby
  - start passive holddown
*Dec  2 16:11:31.855: HSRP: Gi0/1 Grp 1 Standby router is unknown, was
  local
*Dec  2 16:11:31.855: HSRP: Gi0/1 Grp 1 Standby -> Active
```




Use **debug standby terse** to view the HSRP events as R1 is powered down and R2 assumes the role of active HSRP router for the 172.16.10.0/24 network.

## HSRP Debug Commands (Cont.)

R1 Initiates Coup to Become Active HSRP Router

```
R1#
*Dec 2 18:01:30.183: HSRP: Gi0/1 Nbr 172.16.10.2 Adv in, active 0
passive 1
*Dec 2 18:01:30.183: HSRP: Gi0/1 Nbr 172.16.10.2 created
*Dec 2 18:01:30.183: HSRP: Gi0/1 Nbr 172.16.10.2 is passive
*Dec 2 18:01:32.443: HSRP: Gi0/1 Nbr 172.16.10.2 Adv in, active 1
passive 1
*Dec 2 18:01:32.443: HSRP: Gi0/1 Nbr 172.16.10.2 is no longer passive
*Dec 2 18:01:32.443: HSRP: Gi0/1 Nbr 172.16.10.2 destroyed
*Dec 2 18:01:32.443: HSRP: Gi0/1 Grp 1 Coup in 172.16.10.2 Listen
pri 150 vIP 172.16.10.1
*Dec 2 18:01:32.443: HSRP: Gi0/1 Grp 1 Active: j/Coup rcvd from higher
pri router (150/172.16.10.2)
*Dec 2 18:01:32.443: HSRP: Gi0/1 Grp 1 Active router is 172.16.10.2,
was local
*Dec 2 18:01:32.443: HSRP: Gi0/1 Nbr 172.16.10.2 created
*Dec 2 18:01:32.443: HSRP: Gi0/1 Nbr 172.16.10.2 active for group 1
*Dec 2 18:01:32.443: HSRP: Gi0/1 Grp 1 Active -> Speak
*Dec 2 18:01:32.443: %HSRP-5-STATECHANGE: GigabitEthernet0/1 Grp 1
state Active -> Speak
*Dec 2 18:01:32.443: HSRP: Gi0/1 Grp 1 Redundancy "hsrp-Gi0/1-1"
state Active -> Speak
*Dec 2 18:01:32.443: HSRP: Gi0/1 Grp 1 Removed 172.16.10.1 from ARP
*Dec 2 18:01:32.443: HSRP: Gi0/1 IP Redundancy "hsrp-Gi0/1-1" update,
Active -> Speak
*Dec 2 18:01:43.771: HSRP: Gi0/1 Grp 1 Speak: d/Standby timer expired
(unknown)
*Dec 2 18:01:43.771: HSRP: Gi0/1 Grp 1 Standby router is local
*Dec 2 18:01:43.771: HSRP: Gi0/1 Grp 1 Speak -> Standby
```



Because R1 is configured with the **standby 1 preempt** command, it initiates a coup and assumes the role of active router. R2 actively listens to hello messages during the Speak state until it confirms that R1 is the new active router and R2 is the new standby router.

## HSRP Debug Commands (Cont.)

R1 is Administratively Shutdown and Resigns as Active HSRP Router

```
R1(config)# interface g0/1
R1(config-if)# shutdown
R1(config-if)#
*Dec 2 17:36:20.275: %HSRP-5-STATECHANGE: GigabitEthernet0/1
  Grp 1 state Active -> Init
*Dec 2 17:36:22.275: %LINK-5-CHANGED: Interface
  GigabitEthernet0/1, changed state to administratively down
*Dec 2 17:36:23.275: %LINEPROTO-5-UPDOWN: Line protocol on
  Interface GigabitEthernet0/1, changed state to down
R1(config-if)#
!-----
R2#
*Dec 2 17:36:30.699: HSRP: Gi0/1 Grp 1 Resign in 172.16.10.2
  Active pri 150 vIP 172.16.10.1
*Dec 2 17:36:30.699: HSRP: Gi0/1 Grp 1 Standby: i/Resign rcvd
  (150/172.16.10.2)
*Dec 2 17:36:30.699: HSRP: Gi0/1 Grp 1 Active router is local,
  was 172.16.10.2
*Dec 2 17:36:30.699: HSRP: Gi0/1 Nbr 172.16.10.2 no longer
  active for group 1 (Standby)
*Dec 2 17:36:30.699: HSRP: Gi0/1 Nbr 172.16.10.2 Was active
  or standby - start passive holddown
*Dec 2 17:36:30.699: HSRP: Gi0/1 Grp 1 Standby router is
  unknown, was local
*Dec 2 17:36:30.699: HSRP: Gi0/1 Grp 1 Standby -> Active
```

