

Chapter 6: First-Hop Redundancy



CCNP SWITCH: Implementing Cisco IP Switched Networks


Cisco | Networking Academy®
Mind Wide Open™



Chapter 6 Objectives

- Overview of FHRP and HSRP
- Configure and verify VRRP
- Configure and verify GLBP

Overview of FHRP and HSRP



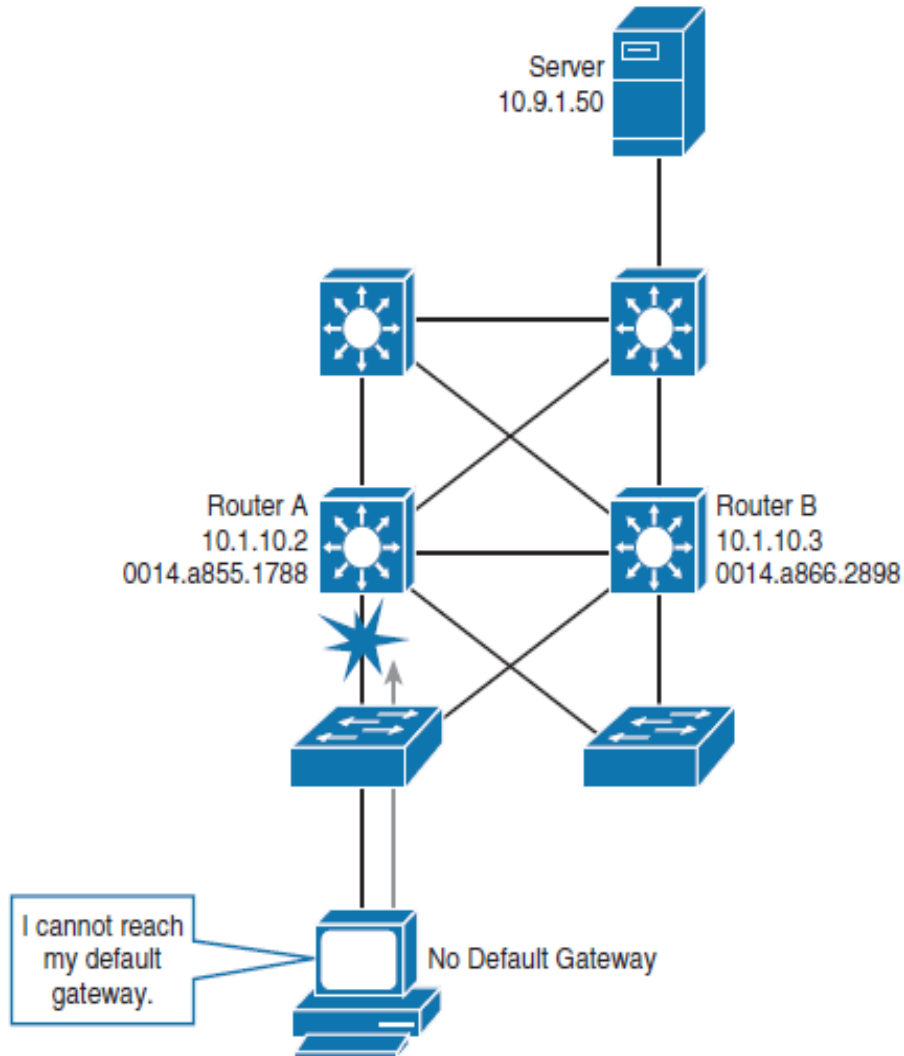


Overview of FHRP and HSRP

- The need for first-hop redundancy protocols
- HSRP overview
- HSRP state transitions
- Aligning HSRP with STP topology
- Configuring and tuning HSRP
- Load sharing with HSRP
- Options HSRP has for tracking
- Configuring HSRP interface tracking
- Configuring object tracking in combination with HSRP
- Configuring HSRP authentication
- Tuning HSRP timers
- The differences between HSRP Versions 1 and 2



The Need for First-Hop Redundancy

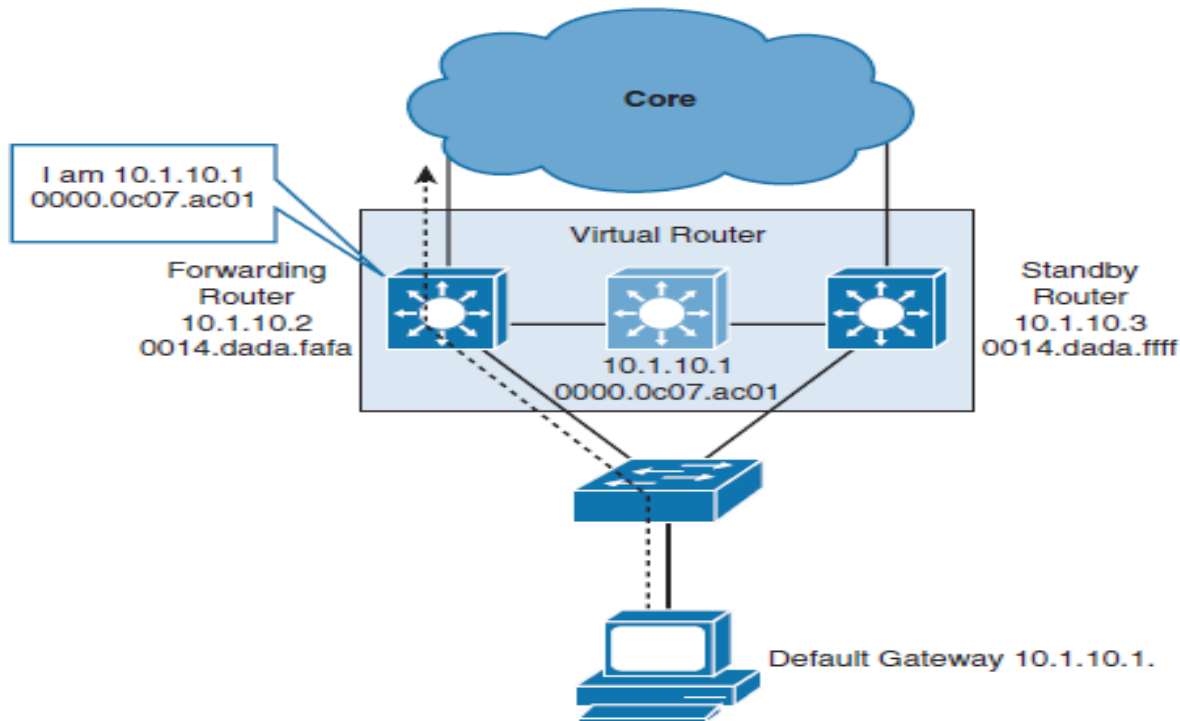


- Network hosts are configured with a **single default gateway** IP address
- If the router whose IP address serves as the default gateway to the network host fails, a network host will be unable to send packets to another subnet



The Need for First-Hop Redundancy

- With first-hop router redundancy, a **set of routers** or **Layer 3 switches** work together to present the **illusion of a single virtual router** to the hosts on the LAN.
- By sharing an IP address and a MAC (Layer 2) address, two or more routers can act as a **single “virtual” router**



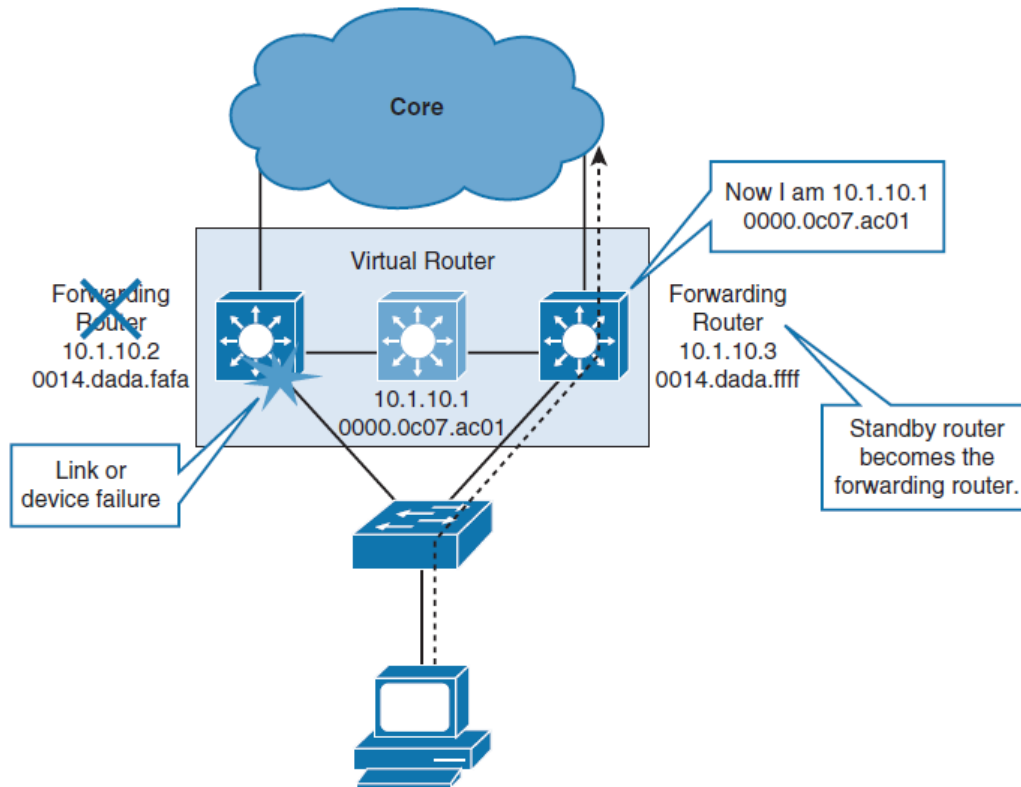


Hot Standby Routing Protocol - HSRP Overview

- When frames are to be sent from the workstation to the default gateway, the workstation uses **ARP to resolve the MAC address** that is associated with the IP address of the default gateway.
- The ARP resolution will return the **MAC address of the virtual router**.
- Frames that are sent to the MAC address of the virtual router can then be physically **processed by an active router** that is part of that virtual router group.
- The physical router that forwards this traffic is **transparent to the network hosts**.
- The redundancy protocol provides the mechanism for determining which router should take the active role in forwarding traffic and determining when that role must be taken over by a standby router.



HSRP Overview



When the forwarding router or a **link to it fails**

- The standby router stops seeing hello messages from the forwarding router.
- The standby router assumes the role of the forwarding router.
- As the new forwarding router assumes both the IP and MAC addresses of the virtual router, the end stations see no disruption in service.



HSRP Overview

- **HSRP active and standby routers** send **hello messages** to multicast address 224.0.0.2 (all routers) for Version 1, or 224.0.0.102 for Version 2, using User Datagram Protocol (UDP) port 1985.
- Hello messages are used to communicate between routers in the HSRP group.
- All the routers in the HSRP group need to be L2 adjacent so that hello packets can be exchanged.



HSRP Router Roles

All the routers in an HSRP group have specific roles and interact in specific manners:

■ Virtual router

- An IP and MAC address pair that end devices have configured as their **default gateway**.
- The **active router processes** all packets and frames sent to the virtual router address.
- The virtual router processes no physical frames.
- There is one virtual router in an HSRP group.

■ Active router

- Within an HSRP group, **one router is elected to be the active router**.
- The **active router physically forwards packets** sent to the MAC address of the virtual router.
- There is one active router in an HSRP group.



HSRP Router Roles

■ Standby router

- **Listens for periodic hello messages.** When the active router fails, the other HSRP routers stop seeing hello messages from the active router.
- The standby router then assumes the role of the active router.
- There is one standby router in an HSRP group.

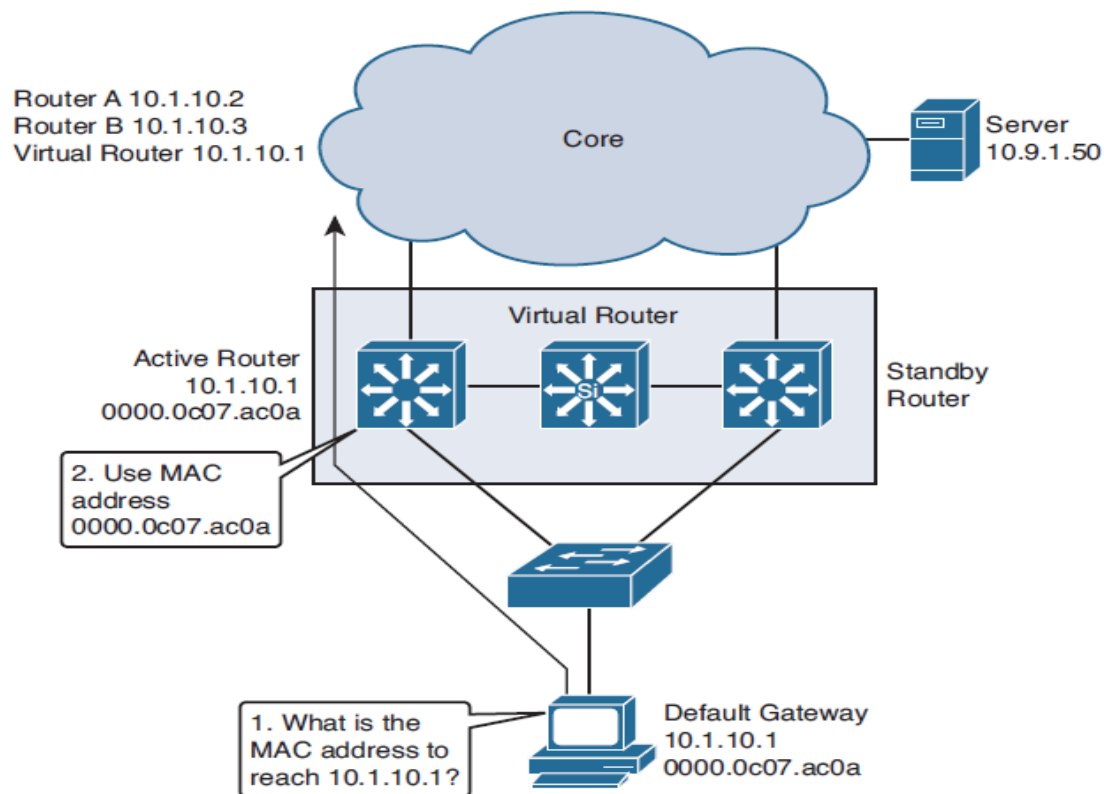
■ Other routers

- There can be more than two routers in an HSRP group, but **only one active** and **one standby router** is possible.
- The other routers **remain in the listen state**, and if both the active and standby routers fail, all routers in the group contend for the active and standby router roles.



HSRP Active Router Operation

- **Router A** assumes the **active role** and forwards all frames addressed to the assigned HSRP MAC address of **0000.0c07.acxx**, where:
 - The first 24 bits will be default CISCO address (i.e. 0000.0c)
 - next 16 bits are HSRP ID (i.e. 07.ac)
 - next 8 bits (xx) will be the group identifier in hexadecimal



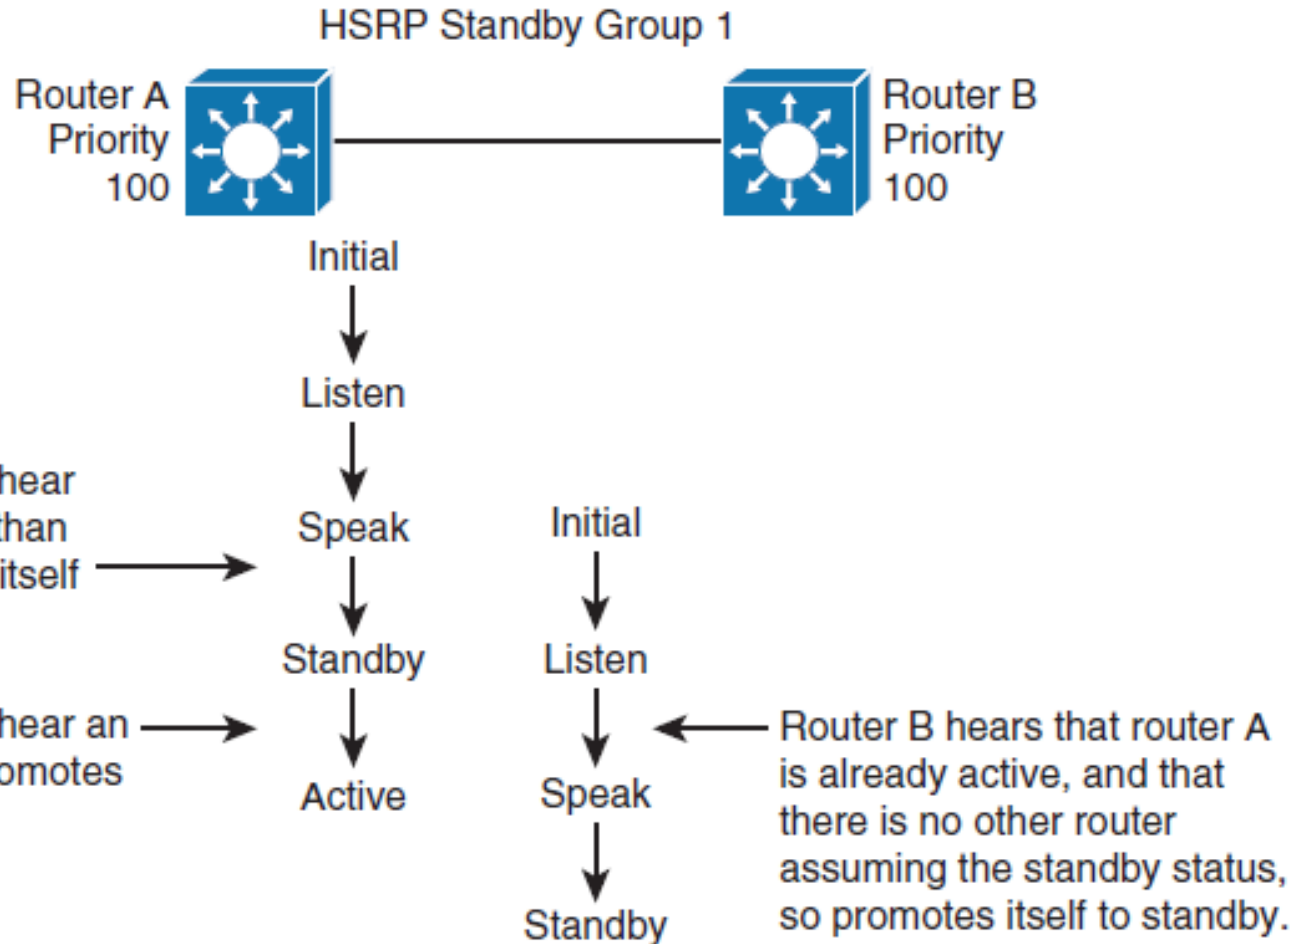


HSRP State Transition

| State | Definition |
|---------|--|
| Initial | The beginning state. The initial state indicates that HSRP does not run. This state is entered via a configuration change or when an interface first comes up. |
| 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 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 the most, one router in the active state in the group. |



HSRP State Transition



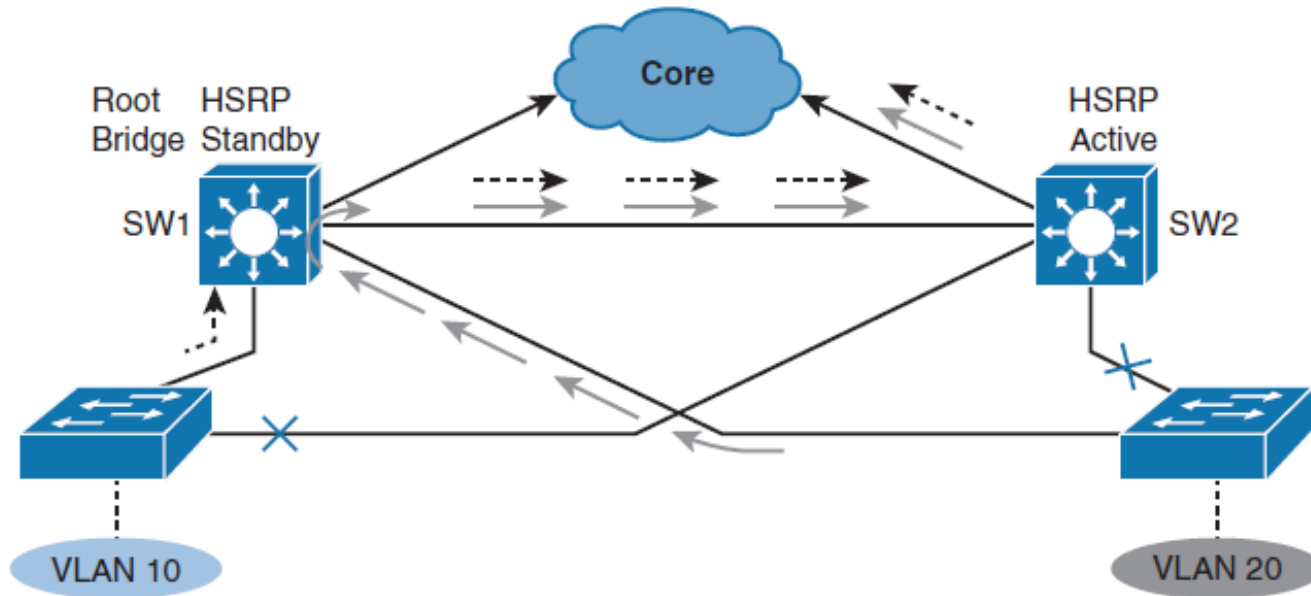


HSRP State Transition

- When two routers participate in an election process, a **priority** can be configured to determine **which router should become active**.
- Without specific priority configuration, each router has a **default priority of 100**, and the router with the **highest IP address** is elected as the active router.
- Regardless of other router priorities or IP addresses, **an active router will stay active by default**.
- **A new election will occur only if the active router is removed**.
- When the standby router is removed, a new election is made to replace the standby router.
- This behavior can change with the **preempt** option.

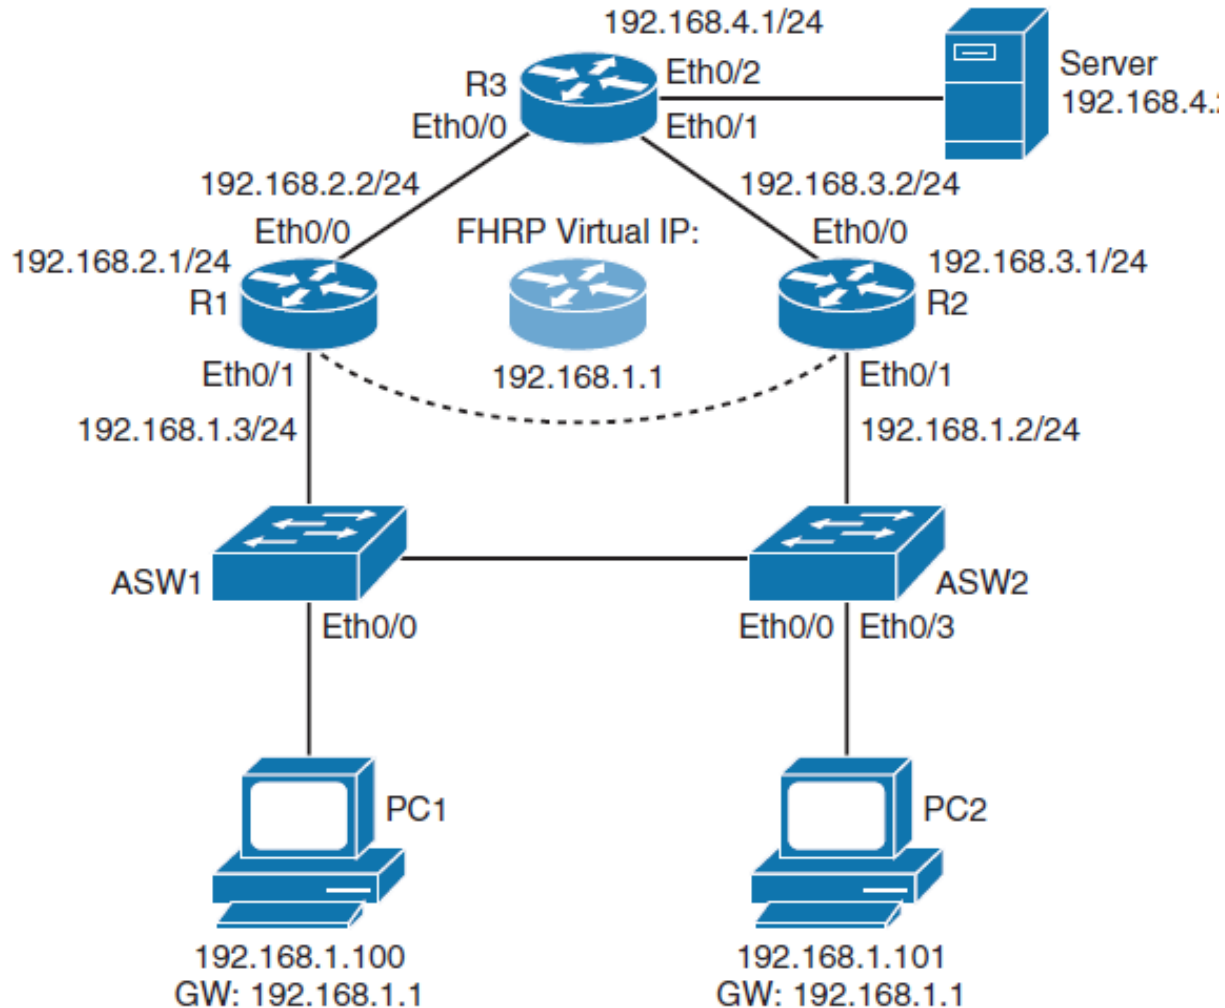


Aligning HSRP with STP Topology



- It is a good practice to configure the **same Layer 3 switch to be both the spanning-tree root and the HSRP active router** for a single VLAN.
- This approach ensures that the Layer 2 forwarding path leads directly to the Layer 3 device that is the HSRP active gateway, thus achieving maximum efficiency.

Configuring and Tuning HSRP





Configuring and Tuning HSRP

Step 1. Configure R1's Ethernet 0/1 (*LAN-facing interface*) with 192.168.1.3/24 IP address and **HSRP standby IP of 192.168.1.1**.

The IP of 192.168.1.1 is HSRP's virtual IP address that is also configured as the default gateway IP address on PC 1 and PC 2:

- R1(config)# **interface ethernet 0/1**
- R1(config-if)# **ip address 192.168.1.3 255.255.255.0**
- R1(config-if)# **standby 1 ip 192.168.1.1**

Step 2. Configure R2's Ethernet 0/1 (*LAN-facing interface*) with 192.168.1.2/24 IP address and **HSRP standby IP of 192.168.1.1**.

Both R1 and R2 must have the same HSRP virtual IP address configured:

- R2(config)# **interface ethernet 0/1**
- R2(config-if)# **ip address 192.168.1.2 255.255.255.0**
- R2(config-if)# **standby 1 ip 192.168.1.1**



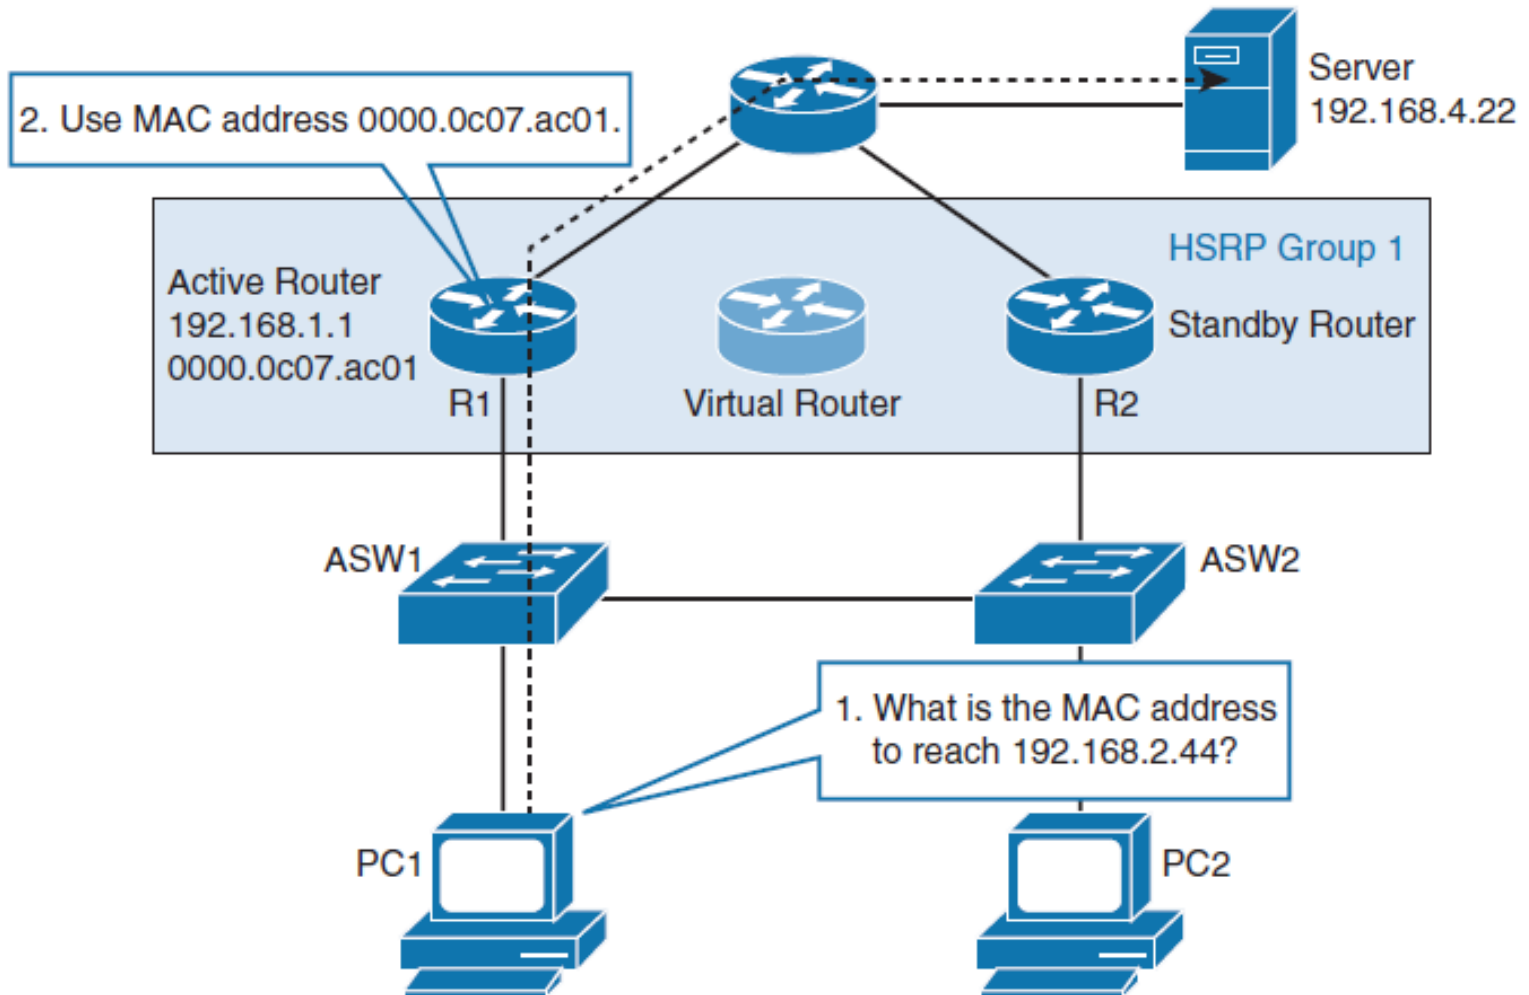
Verify the ARP table

Step 3. On R1, verify the ARP table:

- R1# `show ip arp`

| Protocol | Address | Age (min) | Hardware Addr | Type | Interface |
|----------|-------------|-----------|----------------|------|-------------|
| Internet | 192.168.1.1 | - | 0000.0c07.ac01 | ARPA | Ethernet0/1 |
| Internet | 192.168.1.2 | - | aabb.cc01.ba10 | ARPA | Ethernet0/1 |
| Internet | 192.168.1.3 | 50 | aabb.cc01.bb10 | ARPA | Ethernet0/1 |
| Internet | 192.168.2.1 | - | aabb.cc01.ba00 | ARPA | Ethernet0/0 |
| Internet | 192.168.2.2 | 51 | aabb.cc01.bc00 | ARPA | Ethernet0/0 |

Forwarding Through the Active Router





Make R2 the Active Router

Configure R2's HSRP group 1 with priority of 110:

- R2(config)# **interface ethernet 0/1**
- R2(config-if)# **standby 1 priority 110**

Configure R1's and R2's Ethernet 0/1 HSRP group 1 interfaces with preemption:

- R1(config)# **interface ethernet 0/1**
- R1(config-if)# **standby 1 preempt**
- R2(config)# **interface ethernet 0/1**
- R2(config-if)# **standby 1 preempt**




Make R2 the Active Router

- R2# show standby brief

P indicates configured to preempt.


|

| Interface | Grp | Pri | P | State | Active | Standby | Virtual IP |
|-----------|-----|-----|---|--------|--------|-------------|-------------|
| Et0/1 | 1 | 110 | P | Active | local | 192.168.1.3 | 192.168.1.1 |






Router R2 Failure Scenario

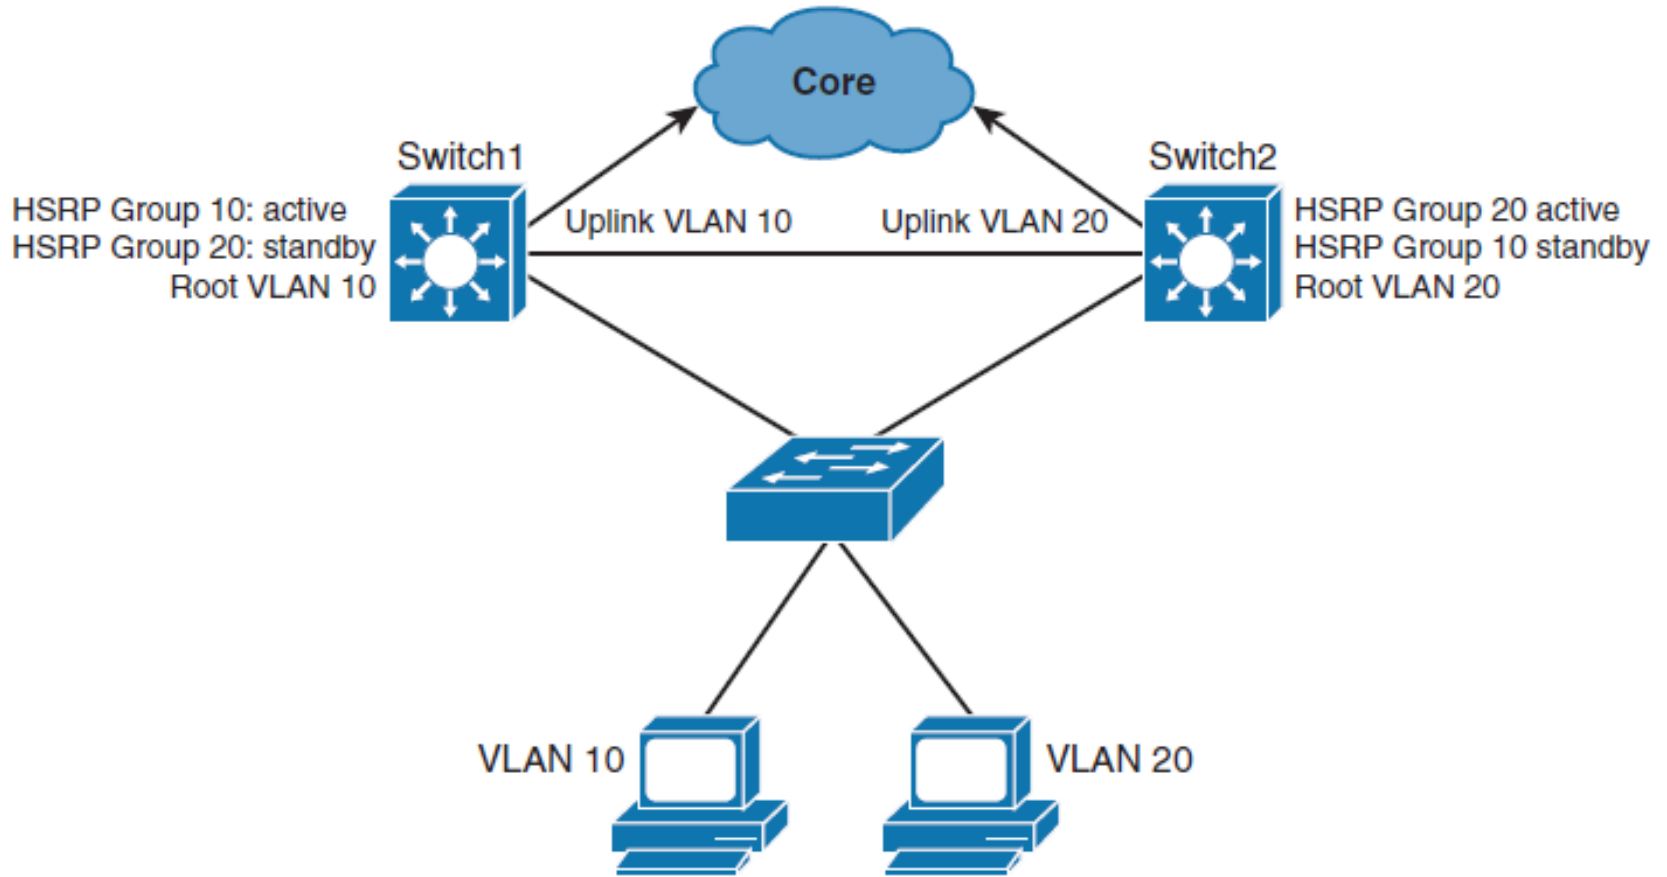


HSRP States After R2 Recover



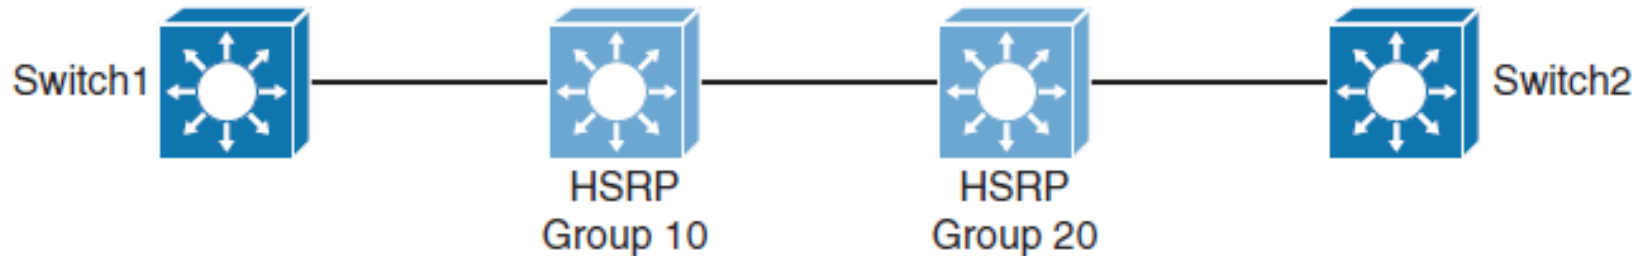


Load Sharing with HSRP






Multigroup HSRP (MHSRP)



```
Switch1(config)# spanning-tree vlan 10 root primary
Switch1(config)# spanning-tree vlan 20 root secondary
Switch1(config)# interface vlan 10
Switch1(config-if)# ip address 10.1.10.2 255.255.255.0
Switch1(config-if)# standby 10 ip 10.1.10.1
Switch1(config-if)# standby 10 priority 110
Switch1(config-if)# standby 10 preempt
Switch1(config-if)# exit
Switch1(config)# interface vlan 20
Switch1(config-if)# ip address 10.1.20.2 255.255.255.0
Switch1(config-if)# standby 20 ip 10.1.20.1
Switch1(config-if)# standby 20 priority 90
Switch1(config-if)# standby 20 preempt
```



Multigroup HSRP (MHSRP)



```
Switch2(config)# spanning-tree vlan 10 root secondary
Switch2(config)# spanning-tree vlan 20 root primary
Switch2(config)# interface vlan 10
Switch2(config-if)# ip address 10.1.10.3 255.255.255.0
Switch2(config-if)# standby 10 ip 10.1.10.1
Switch2(config-if)# standby 10 priority 90
Switch2(config-if)# standby 10 preempt
Switch2(config-if)# exit
Switch2(config)# interface vlan 20
Switch2(config-if)# ip address 10.1.20.3 255.255.255.0
Switch2(config-if)# standby 20 ip 10.1.20.1
Switch2(config-if)# standby 20 priority 110
Switch2(config-if)# standby 20 preempt
```

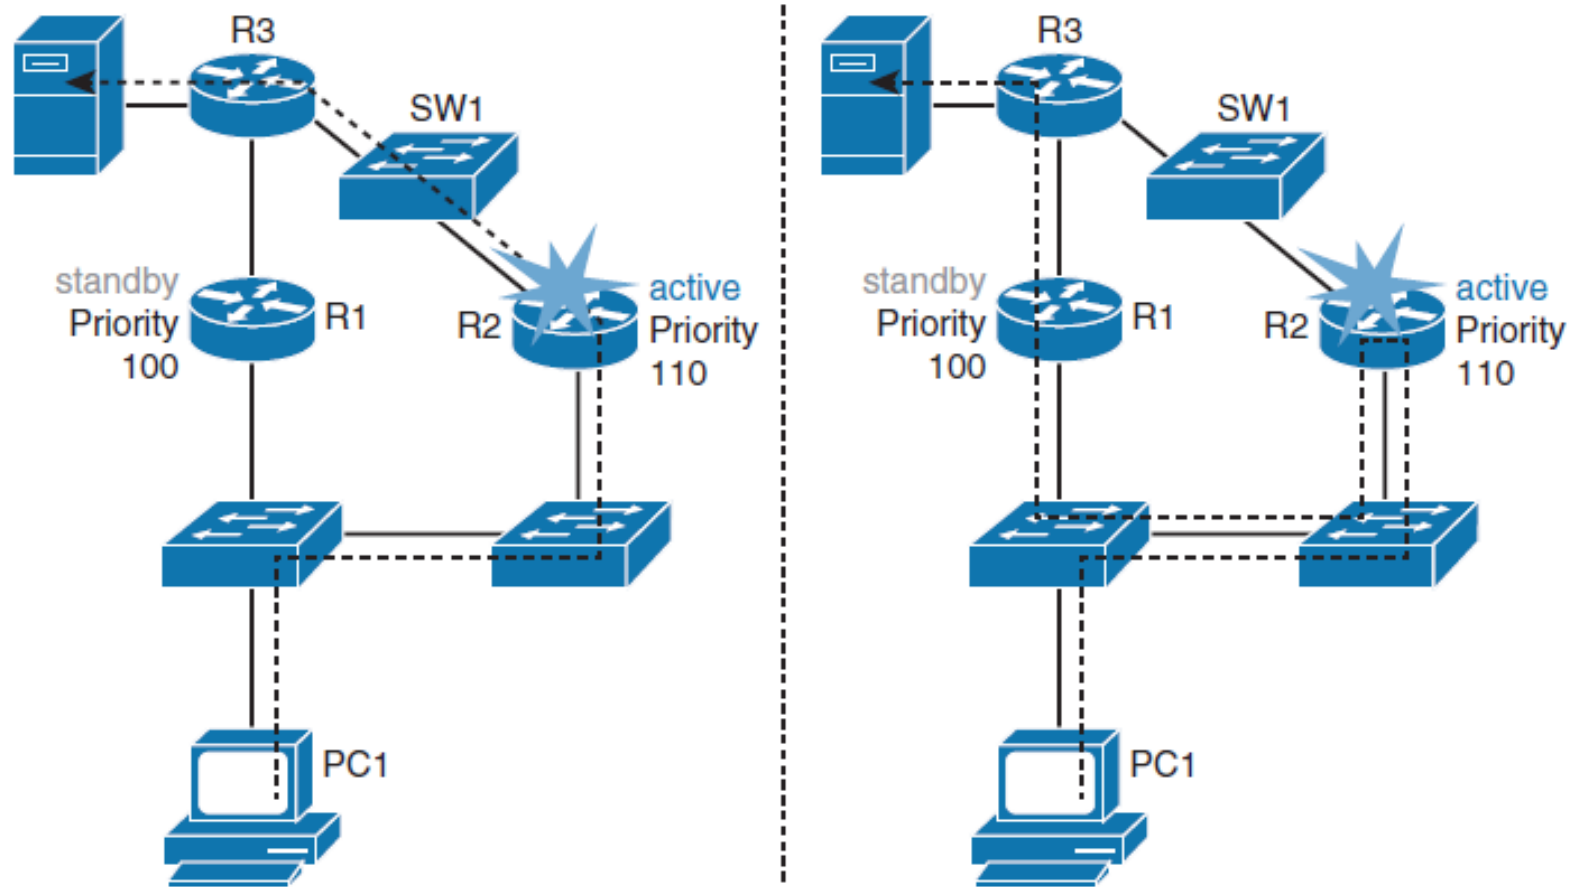


The Need for Interface Tracking with HSRP

- HSRP can **track interfaces or objects** and **decrement priority** if an interface or object fails.
- Interface tracking enables the priority of a standby group router to be automatically adjusted, based on the availability of the router interfaces.
- When a tracked interface becomes unavailable, the HSRP priority of the router is decreased.
- When properly configured, the HSRP tracking feature ensures that a router with an unavailable key interface will relinquish the active router role.
- When the conditions that are defined by the object are fulfilled, the router priority remains the same.
- As soon as the verification that is defined by the object fails, the router priority is decremented.
- The amount of decrease can be configured.
- The default value is 10.

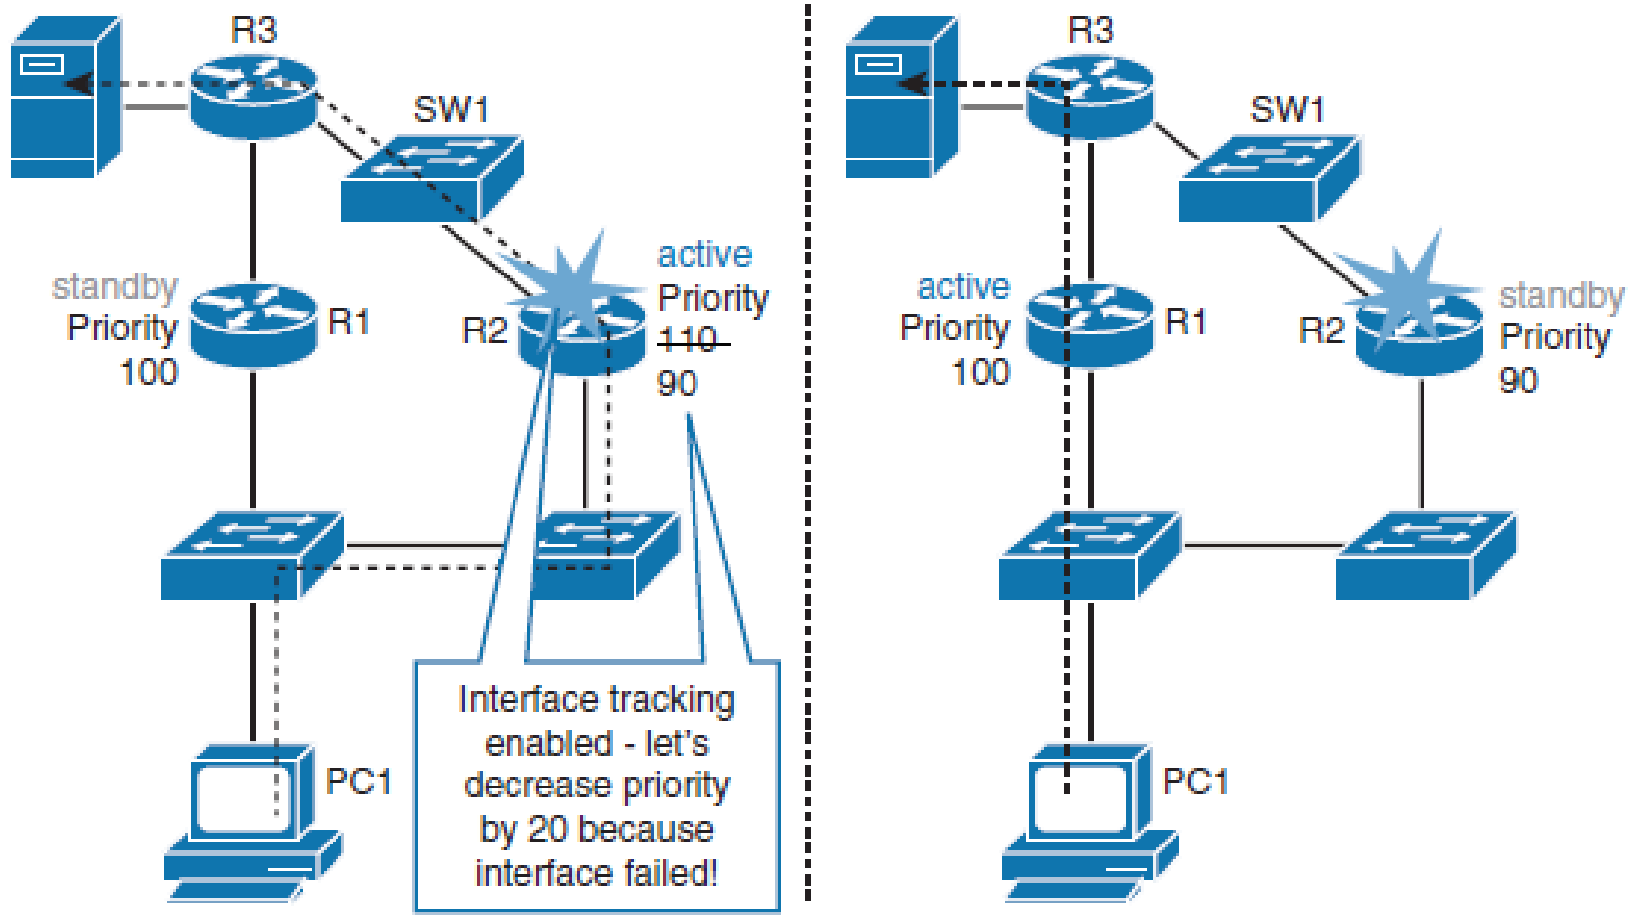


HSRP Interface Tracking



- HSRP has a built-in mechanism for **detecting link failures** and starting the HSRP reelection process.

HSRP with Interface Tracking On





HSRP Tracking Configuration

- R2(config)# **interface ethernet 0/1**
- R2(config-if)# **ip address 192.168.10.2**
- R2(config-if)# **standby 10 ip 192.168.10.1**
- R2(config-if)# **standby 10 priority 110**
- R2(config-if)# **standby 10 preempt**
- R2(config-if)# **standby 10 track ethernet1/1 20**

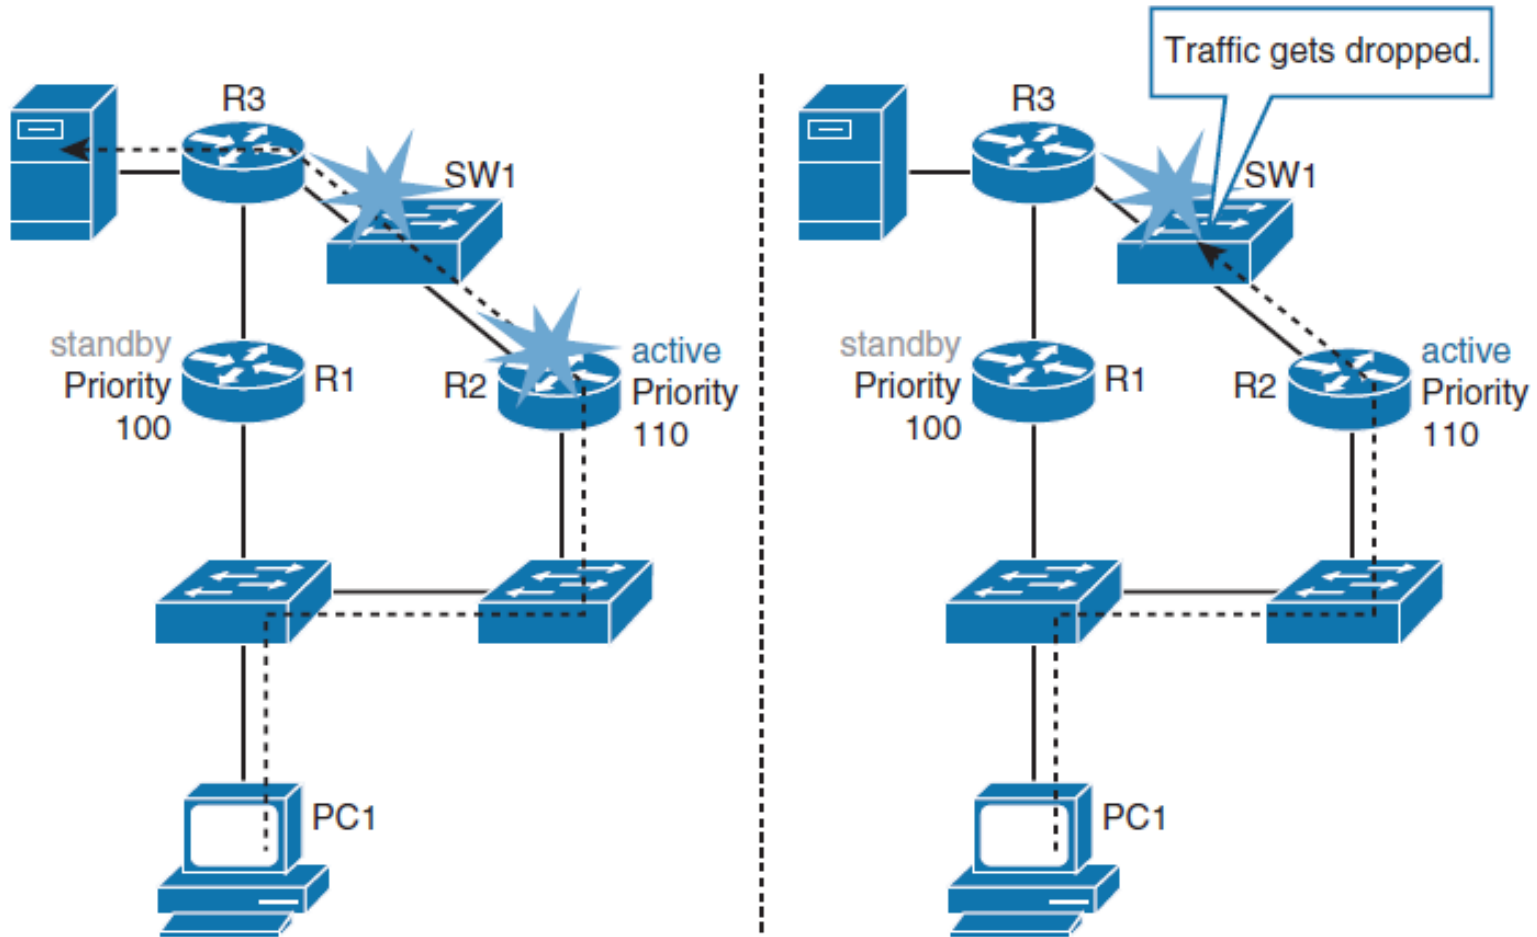


HSRP Tracking Configuration Arguments

| Variable | Description |
|---------------------------|--|
| <i>group-number</i> | (Optional) Indicates the group number on the interface to which the tracking applies. The default number is 0. |
| <i>type</i> | Indicates the interface type (combined with the interface number) that will be tracked. |
| <i>number</i> | Indicates the interface number (combined with the interface type) that will be tracked. |
| <i>interface-priority</i> | (Optional) Indicates the amount by which the hot standby priority for the router is decremented when the interface becomes disabled. The priority of the router is incremented by this amount when the interface becomes available. The default value is 10. |

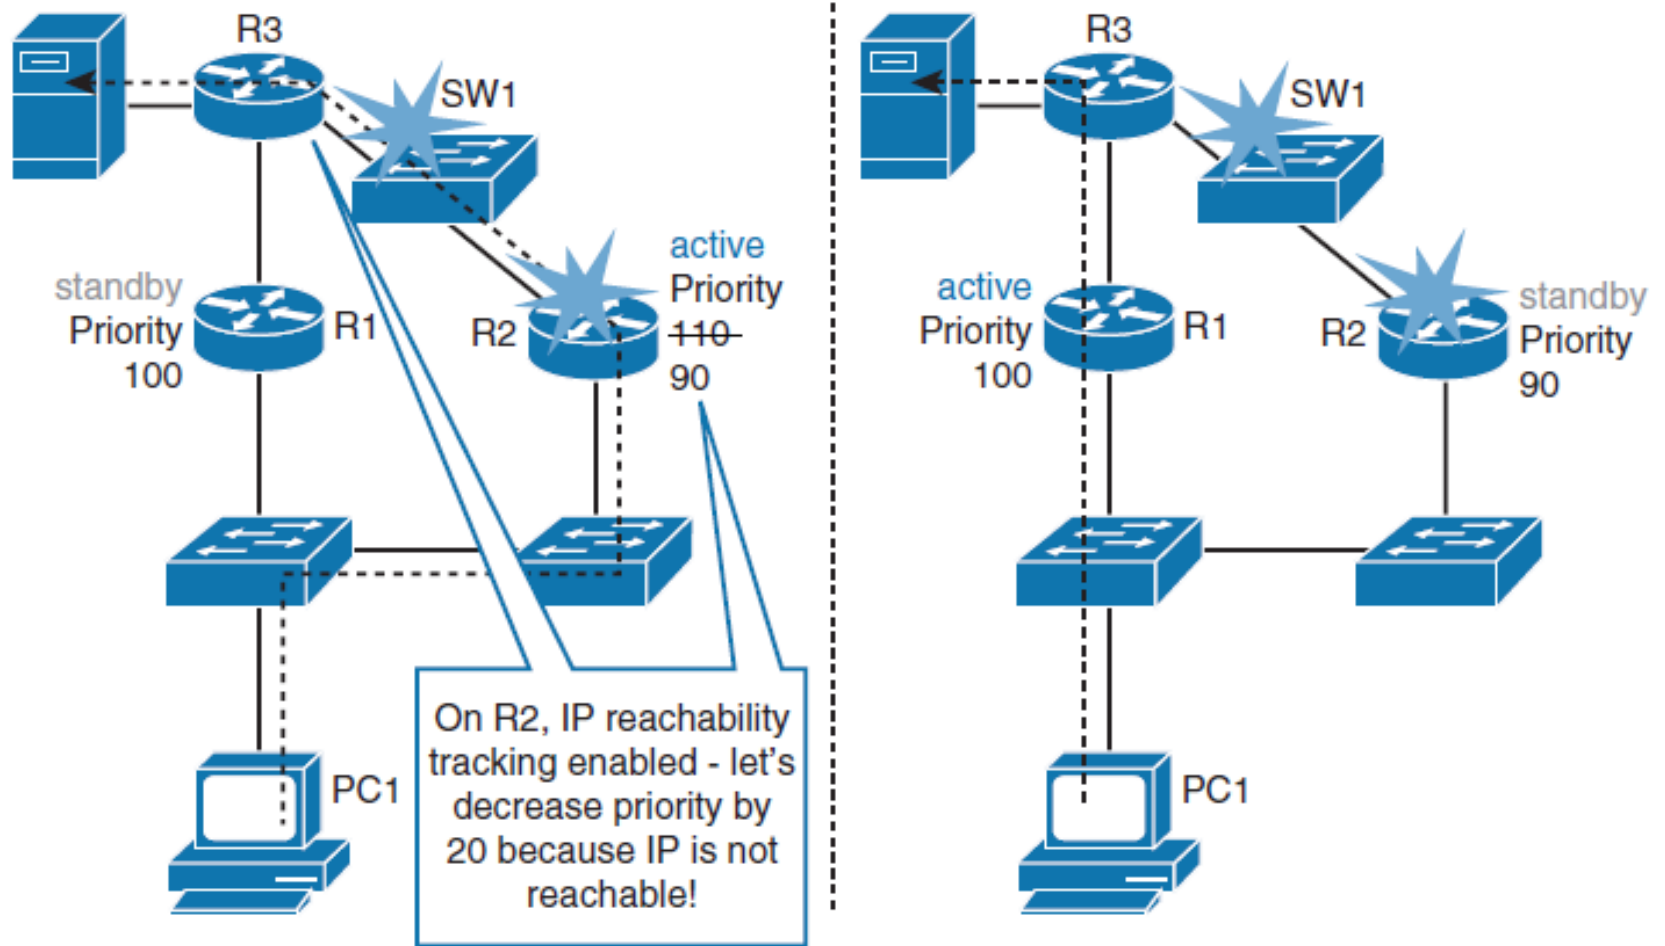


HSRP and Object Tracking





HSRP With Object Tracking





HSRP and Object Tracking Configuration

First, define an IP SLA ICMP echo test:

- `R2(config)# ip sla 10`
- `R2(config-ip-sla)# icmp-echo 192.168.3.2`
- `R2(config-ip-sla-echo)# frequency 5`
- `R2(config-ip-sla-echo)# ip sla schedule 10 life forever start-time now`

Then create an object and track the IP SLA instance:

- `R2(config)# track 100 ip sla 10`

Then configure HSRP to track an object and decrement priority if the test fails:

- `R2(config)# interface ethernet 0/1`
- `R2(config-if)# standby 1 track 100 decrement 20`



Tracked objects

- Tracked objects are defined in global configuration with the keyword **track** , followed by an object number.
- Although IP SLA is just one of the options that can be tracked, as shown in the following syntax, you can track up to 500 objects:
 - `Switch(config)# track 1 ?`
 - **interface** Select an interface to track
 - **ip** IP protocol
 - **list** Group objects in a list



Tracked objects

- Tracked objects offer a vast group of possibilities.
- A few options that are commonly available include the following
 - **An interface**
 - This performs a similar function like the HSRP interface tracking mechanism, but with advanced features. This tracking object can not only verify the interface status (line protocol) but also whether **IP routing is enabled**, whether an **IP address is configured** on the interface, and whether the **interface state is up**, before reporting to the tracking client that the interface is up.
 - **IP route**
 - A tracked IP-route object is considered up and reachable **when a routing table entry exists** for the route and the route is accessible. To provide a common interface to tracking clients, route metric values are normalized to the range of 0 to 255, where 0 is connected and 255 is inaccessible. You can track **route reachability**, or even **metric values**, to determine best-path values to the target network.



Tracked objects (continue)

■ IP SLA

- This special case allows you to track advanced parameters such as **IP reachability**, **delay**, or **jitter**.

■ A list of objects

- You can track several objects and interrelate their results to determine whether one or several of them should trigger the “success” or “fail” condition.



Configuring HSRP Authentication

- HSRP provides the following two types of authentication:
 - Plain text
 - Message digest 5 (MD5) algorithm

To configure **plain-text authentication**, use the following interface configuration command on HSRP peers:

- `Switch(config-if)# standby group authentication string`

To configure **MD5 authentication**, use the following interface configuration command on HSRP peers:

- `Switch(config-if)# standby group authentication md5 key-string [0 | 7] string`



Configuring HSRP Authentication

- To configure MD5 authentication using key chains, use the following command sequence:
- Switch(config)# **key chain** *chain-name*
- Switch(config-keychain)# **key** *key-number*
- Switch(config-keychain-key)# **key-string** [0 | 7] *string*
- Switch(config-keychain-key)# **exit**
- Switch(config)# **interface** *interface-slot/number*
- Switch(config-if)# **standby group authentication md5 key-chain** *chain-name*



Tuning HSRP Timers

- By default, the HSRP **hello time is 3 seconds**, and the **hold time is 10 seconds**, which means that the failover time could be as much as 10 seconds for clients to start communicating with the new default gateway.
- In some cases, this interval may be excessive for application support.
- The hello-time and the hold-time parameters are configurable.
- To configure the time between the hello messages and the time before other group routers declare the active or standby router to be non-functioning, enter this command in the interface configuration mode:
- `Switch(config-if)# standby [group-number] timers [msec] hellotime [msec] holdtime`



Preemption Delay

- Preemption is an important feature of HSRP that **allows the primary router to resume the active role** when it comes back online after a failure or a maintenance event.
- Preemption is a desired behavior as it forces a predictable routing path for the VLAN traffic during normal operations.
- It also ensures that the Layer 3 forwarding path for a VLAN parallels the Layer 2 STP forwarding path whenever possible.
- When a preempting device is rebooted, HSRP preemption communication should not begin until the distribution switch has established full connectivity to the rest of the network.
- This situation allows the routing protocol convergence to occur more quickly, after the preferred router is in an active state.
- To accomplish this, measure the system boot time and set the HSRP preemption delay to a value that is about 50 percent greater than device's boot time



Configuring HSRP Preemption Delay Timers

```

switch(config)# interface vlan 10
switch(config-if)# ip address 10.1.1.2 255.255.255.0
switch(config-if)# standby 10 ip 10.1.1.1
switch(config-if)# standby 10 priority 110
switch(config-if)# standby 10 preempt
switch(config-if)# standby 10 timers msec 200 msec 750
switch(config-if)# standby 10 preempt delay minimum 225
  
```

- For example, if the boot time for the distribution device is 150 seconds, the preempt delay should be set to 225 seconds



HSRP Versions

There are two HSRP versions available on most Cisco routers and Layer 3 switches:

- HSRPv1 and HSRPv2.
- Version 1 is a default version on Cisco IOS devices.
- **HSRPv2 allows group numbers up to 4095, thus allowing you to use VLAN number as the group number.**
- HSRP Version 2 must be enabled on an interface before HSRP IPv6 can be configured.
- **HSRP Version 2 will not interoperate with HSRP Version 1.**
- All devices in an HSRP group must have the same version configured; otherwise, the hello messages are not understood.



HSRP Versions

- An interface **cannot operate both Version 1 and Version 2** because they are mutually exclusive.
- The MAC address of the virtual router and the multicast address for the hello messages are different with Version 2.
- HSRPv2 uses the new IP multicast address 224.0.0.102 to send the hello packets instead of the multicast address of 224.0.0.2, which is used by Version 1.
- To enable HSRP Version 2, enable the following configuration:
 - `Switch(config-if) standby hsrp-number version 2`

Configuring Layer 3 Redundancy with VRRP - Virtual Router Redundancy Protocol





Configuring Layer 3 Redundancy with VRRP

Upon completing this section, you will be able to do the following:

- Describe the idea behind VRRP
- Configure and verify VRRP
- Describe the differences between HSRP and VRRP
- Describe tracking options with VRRP
- Configure VRRP interface object tracking



About VRRP

- VRRP is an open standard **alternative to HSRP**.
- VRRP is similar to HSRP, both in operation and configuration.
- The **VRRP master** is analogous to the **HSRP active** gateway, and the **VRRP backup** is analogous to the **HSRP standby** gateway.
- A VRRP group has one master device and one or multiple backup devices.
- A device with the highest priority is the elected master. Priority can be a number between 0 and 255.
 - Priority value 0 has a special meaning; it indicates that the current master has stopped participating in VRRP.
 - This setting is used to trigger backup devices to quickly transition to master without having to wait for the current master to time out.

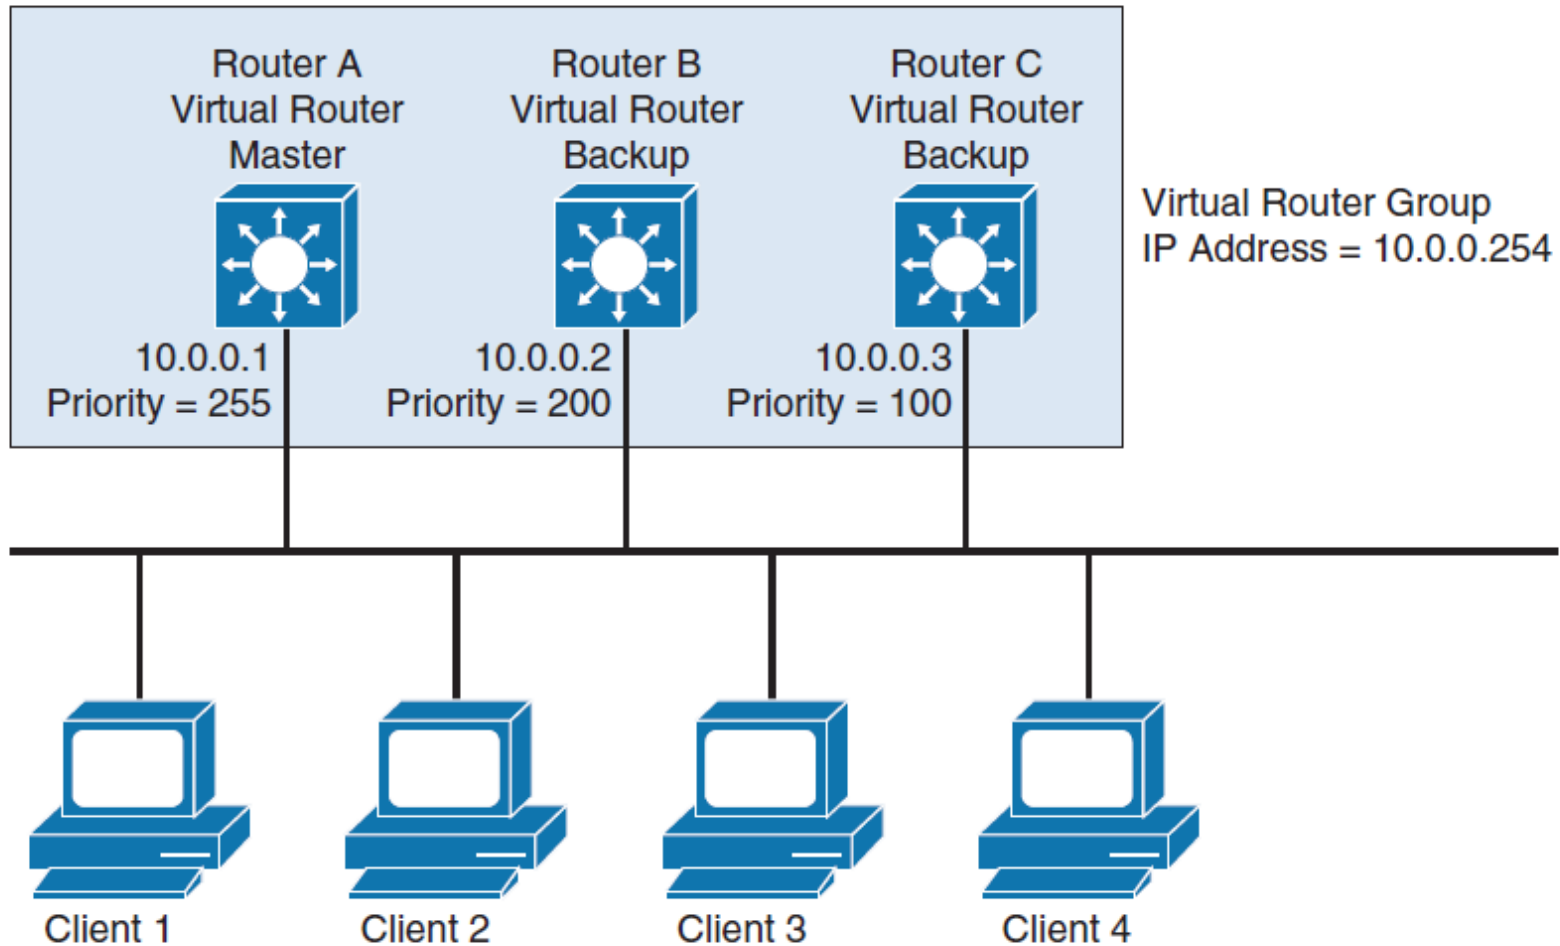


About VRRP

- VRRP differs from HSRP in that it allows you to use an address of one of the physical VRRP group members as a virtual IP address.
 - In this case, the device with the used physical address is a VRRP master whenever it is available.
- The master is the only device that sends advertisements (analogous to HSRP hellos).
- Advertisements are sent to the 224.0.0.18 multicast address, protocol number 112.
- The default advertisement interval is 1 second. The default hold time is 3 seconds.
- HSRP, in comparison, has the default hello timer set to 3 seconds and the hold timer to 10 seconds.
- Like with HSRP, load sharing is also available with VRRP. Multiple virtual router groups can be configured
- Contrary to HSRP, preemption is enabled by default with VRRP.



About VRRP






HSRP and VRRP Differences

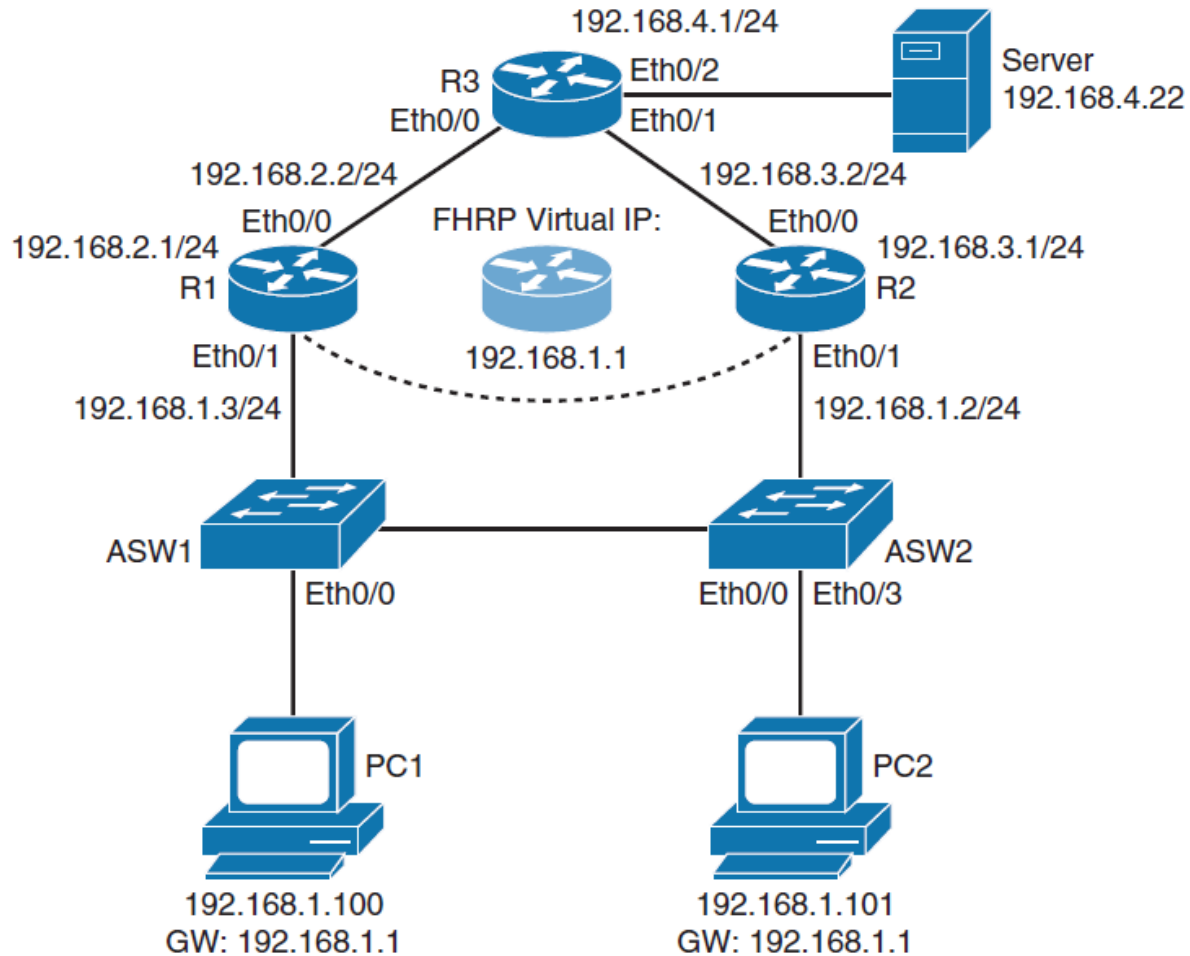
| HSRP | VRRP |
|---|--|
| Cisco proprietary. | Industry standard. |
| 1 active, 1 standby, several candidates. | 1 master, several backups. |
| Virtual IP is different from real IP addresses. | Virtual IP address can be the same as the real IP address of one of the group members. |
| Uses 224.0.0.2. | Uses 224.0.0.18. |
| Can track interfaces or objects. | Can track only objects. |
| Default timers: hello 3 sec; hold time 10 sec. | Default timers: hello 1 sec; hold time 3 sec. |
| Authentication supported. | Authentication no longer in RFC, but still supported in Cisco IOS. |



Configuring VRRP and Spotting the Differences from HSRP



IP Addressing for the VRRP Configuration





Configuring VRRP

Step 1. Configure R1's Ethernet 0/1 with IP address 192.168.1.3 and VRRP virtual IP address 192.168.1.1:

- R1(config)# **interface ethernet 0/1**
- R1(config-if)# **ip address 192.168.1.3 255.255.255.0**
- R1(config-if)# **vrrp 1 ip 192.168.1.1**

Configure R2's Ethernet 0/1 with IP address of 192.168.1.2 and VRRP virtual IP address of 192.168.1.1:

- R2(config)# **interface ethernet 0/1**
- R2(config-if)# **ip address 192.168.1.2 255.255.255.0**
- R2(config-if)# **vrrp 1 ip 192.168.1.1**

Step 2. Configure R2's Ethernet 0/1 with VRRP priority of 110:

- R2(config-if)# **vrrp 1 priority 110**



Verify the VRRP Status

```
R1# show vrrp
```

```
Ethernet0/1 - Group 1
```

```
State is Backup
```

```
Virtual IP address is 192.168.1.1
```

```
Virtual MAC address is 0000.5e00.0101
```

```
Advertisement interval is 1.000 sec
```

```
Preemption enabled
```

```
Priority is 100
```

```
Master Router is 192.168.1.2, priority is 110
```

```
Master Advertisement interval is 1.000 sec
```

```
Master Down interval is 3.609 sec (expires in 3.049 sec)
```

```
R2# show vrrp brief
```

| Interface | Grp | Pri | Time | Own | Pre | State | Master addr | Group addr |
|-----------|-----|-----|------|-----|-----|--------|-------------|-------------|
| Et0/1 | 1 | 110 | 3570 | | Y | Master | 192.168.1.2 | 192.168.1.1 |



VRRP and Authentication

Configure MD5 authentication for VRRP on R1's Ethernet 0/1 interface:

- R1(config)# **interface ethernet 0/1**
- R1(config-if)# **vrrp 1 authentication md5 key-string MyVRRP**
- %VRRP-4-BADAUTHTYPE: Bad authentication from 192.168.1.2, group 1, type 0, expected 254.

Configure MD5 authentication for VRRP on R2's Ethernet 0/1 interface:

- R2(config)# **interface ethernet 0/1**
- R2(config-if)# **vrrp 1 authentication md5 key-string MyVRRP**

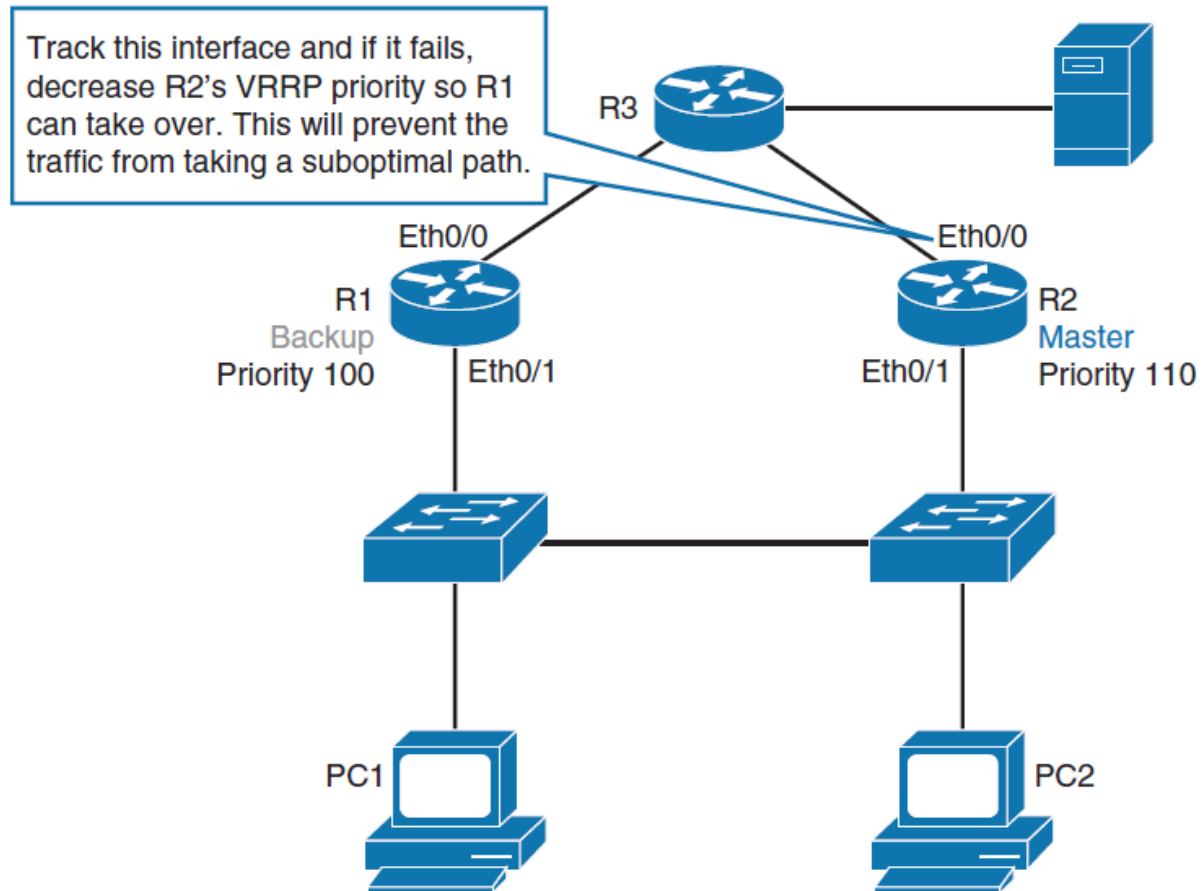
R1's CLI:

- %VRRP-6-STATECHANGE: Et0/1 Grp 1 state Master -> Backup



Tracking and VRRP

- VRRP **does not have a native interface tracking** mechanism, but it does have the ability to track objects.





Tracking and VRRP Configuration

Create a tracked object, where the status of the uplink interface is tracked:

- `R2(config)# track 1 interface ethernet 0/0 line-protocol`

Configure VRRP to track previously created object and decrease VRRP priority by 20 should the uplink fail:

- `R2(config)# interface ethernet 0/1`
- `R2(config-if)# vrrp 1 track 1 decrement 20`

Configuring Layer 3 Redundancy with GLBP - Gateway Load Balancing Protocol





Configuring Layer 3 Redundancy with GLBP

Upon completing this section, you will be able to do the following:

- Describe the basic idea behind GLBP
- Compare GLBP to HSRP
- Describe the possible states of GLBP virtual gateway and virtual forwarder
- Configure and verify GLBP
- Understand GLBP operations
- List and describe GLBP load-balancing options
- Secure GLBP using authentication
- Describe GLBP behavior in VLANs with running STP
- Describe the system of weights and decrements in GLBP



Introducing GLBP

- GLBP shares some concepts with VRRP and HSRP, but the terminology differs, and its **behavior is more dynamic and robust**.
- Although HSRP and VRRP provide gateway resiliency **only the active router within the group forwards the traffic** for the virtual MAC.
- HSRP and VRRP can accomplish load sharing by manually specifying **multiple groups** and assigning multiple default gateways.
- **GLBP is a Cisco proprietary solution** that allows for **automatic selection and simultaneous use of multiple available gateways**, in addition to automatic failover between those gateways.
- **Multiple routers share the load of packets that, from a client's perspective, are sent to a single default gateway address.**
- There is also no need to configure a specific gateway address on an individual host. All hosts can use the same default gateway.



GLBP Roles

- GLBP routers are divided into two roles: a **gateway** and a **forwarder**:
- **GLBP AVG (active virtual gateway)**
 - Members of a GLBP group elect one gateway to be the AVG for that group.
 - Other group members provide a backup for the AVG when the AVG becomes unavailable; these will be in standby state.
 - The AVG assigns a virtual MAC address to each member of the GLBP group.
 - The AVG listens to the ARP requests for the default gateway IP and replies with a MAC address of one of the GLBP group members, thus load sharing traffic among all the group members.
- **GLBP AVF (active virtual forwarder)**
 - Each gateway assumes responsibility for forwarding packets that are sent to the virtual MAC address that is assigned to that gateway by the AVG.
 - These gateways are known as AVFs. There can be **up to four forwarders** within a GLBP group.
 - All other devices will be secondary forwarders, serving as backup if the current AVF fails.
 - Forwarders that are forwarding traffic for a specific virtual MAC are in the active state and are called AVFs. Forwarders that are serving as backups are in the listen state.



Comparing GLPB to HSRP

| HSRP | GLBP |
|---|--|
| Cisco proprietary, 1994. | Cisco proprietary, 2005. |
| 1 active, 1 standby, several candidates. | Active virtual gateway (AVG): 1 active, 1 standby, several candidates. Active virtual forwarder (AVF): Multiple active, several candidates. |
| Virtual IP is different from real IP addresses. | Virtual IP is different from IPs on interfaces. |
| Uses 224.0.0.2 v1, 224.0.0.102 UDP port 1985. | Uses 224.0.0.102 UDP port 3222. |
| Can track interfaces or objects. | Can track only objects. |
| Default timers: hello 3 sec; hold time 10 sec. | Default timers: hello 3 sec; hold time 10 sec. |
| Authentication supported. | Authentication supported. |



GLBP States

| State | Virtual Gateway | Virtual Forwarder |
|----------|-----------------|-------------------|
| Disabled | ✓ | ✓ |
| Initial | ✓ | ✓ |
| Listen | ✓ | ✓ |
| Speak | ✓ | X |
| Standby | ✓ | X |
| Active | ✓ | ✓ |



GLBP States (Gateway)

Following are the possible virtual gateway states:

- **Disabled:** The virtual IP address has not been configured or learned, but there is some GLBP configuration.
- **Initial:** The virtual IP address has been configured or learned, but configuration is not complete. The interface must be operational on Layer 3 and configured to route IP.
- **Listen:** The virtual gateway is receiving hello packets. It is ready to change to speak state if the active or standby virtual gateway becomes unavailable.
- **Speak:** The virtual gateway is trying to become the active or standby virtual gateway.
- **Standby:** This gateway is next in line to be the active virtual gateway.
- **Active:** This gateway is the AVG, and it is responsible for responding to ARP requests for the virtual IP address.



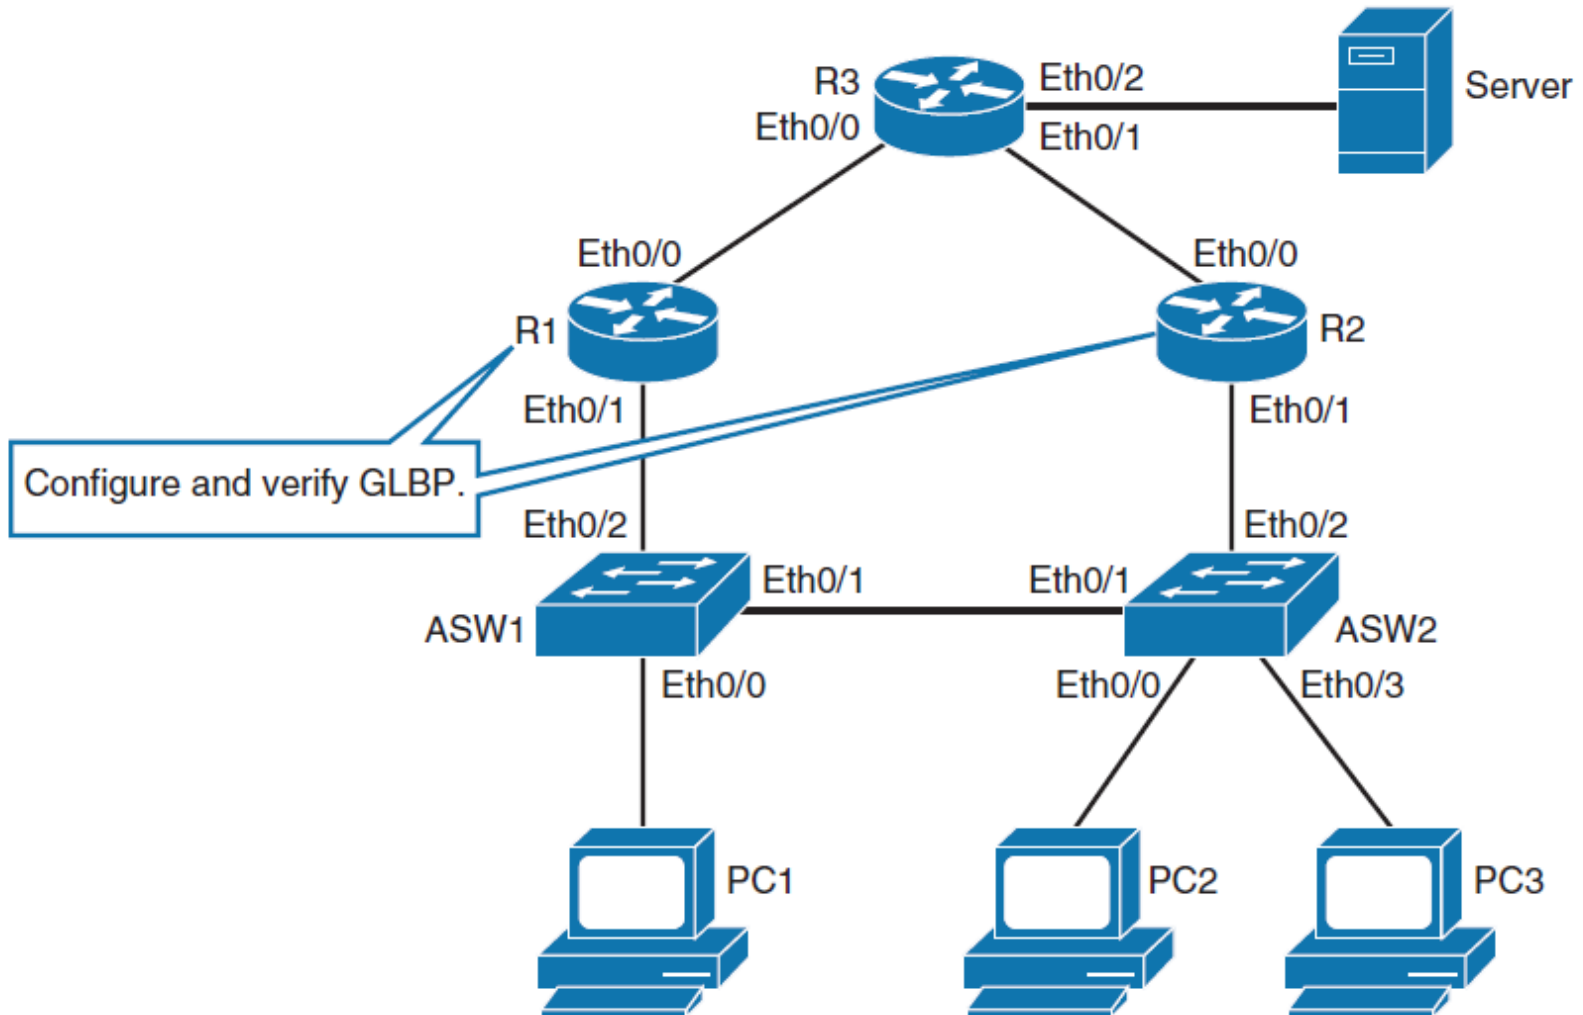
GLBP States (Forwarder)

The following are the possible virtual forwarder states:

- **Disabled:** The virtual MAC address has not been assigned or learned. The disabled virtual forwarder will be deleted shortly. This state is transitory only.
- **Initial:** The virtual MAC address is known but configuration of virtual forwarder is not complete. The interface must be operational on Layer 3 and configured to route IP.
- **Listen:** This virtual forwarder is receiving hello packets and is ready to change to the active state if the active virtual forwarder becomes unavailable.
- **Active:** This gateway is the AVF, and it is responsible for forwarding packets sent to the virtual forwarder's MAC address.

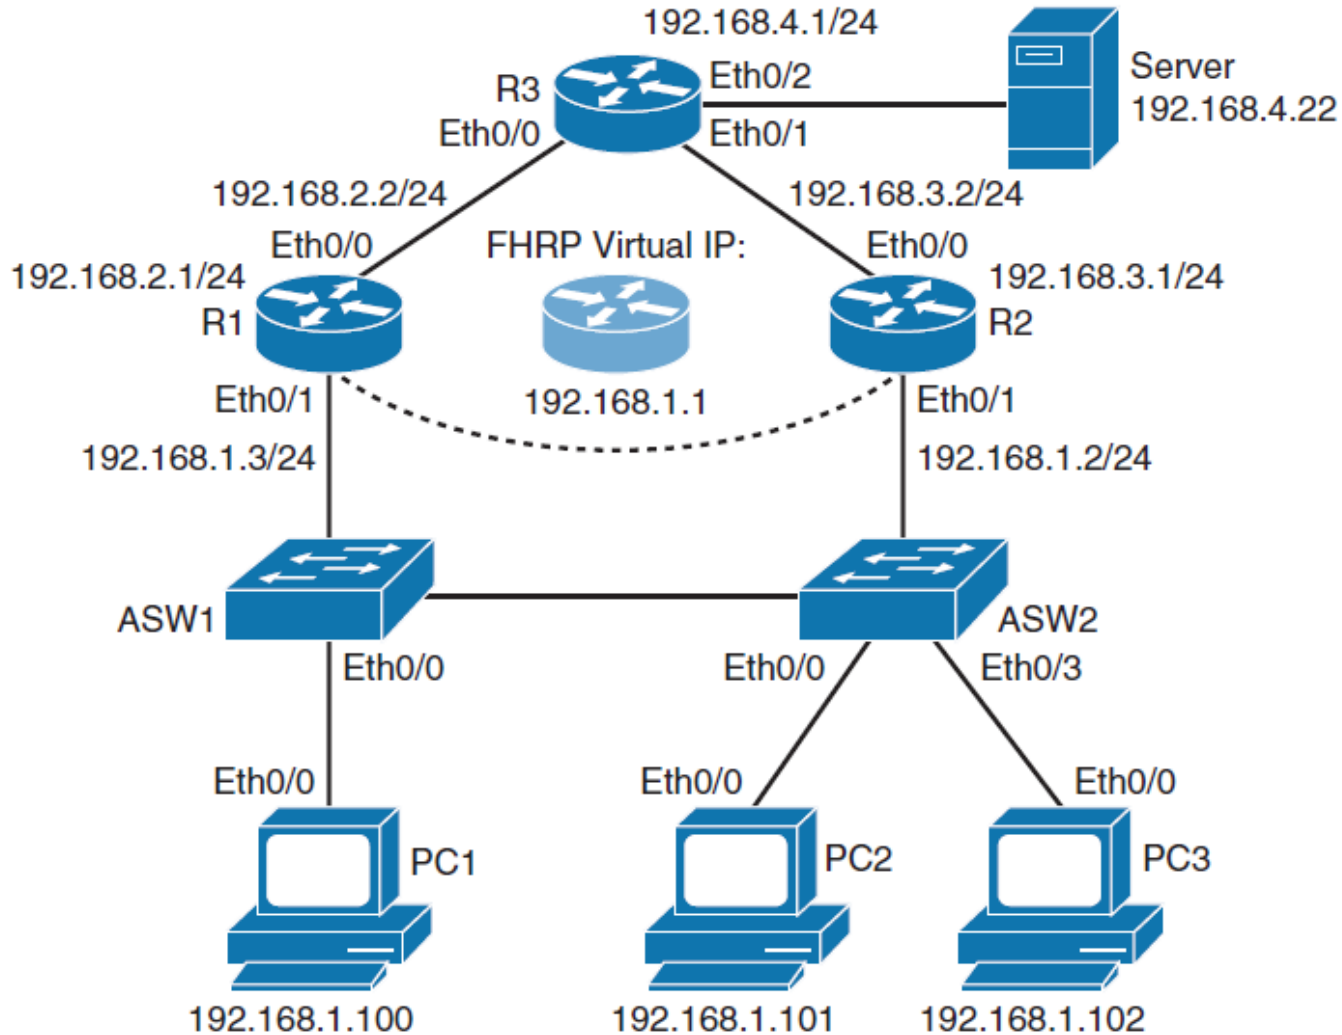


Configuring and Verifying GLBP





IP Addresses Used in GLBP Configuration





GLBP Configuration

Configure R1's Ethernet 0/1 with IP address of 192.168.1.3 and GLBP virtual IP address of 192.168.1.1:

- R1(config)# **interface ethernet 0/1**
- R1(config-if)# **ip address 192.168.1.3 255.255.255.0**
- R1(config-if)# **glbp 1 ip 192.168.1.1**

Configure R2's Ethernet 0/1 with IP address of 192.168.1.2 and GLBP virtual IP address of 192.168.1.1:

- R2(config)# **interface ethernet 0/1**
- R2(config-if)# **ip address 192.168.1.2 255.255.255.0**
- R2(config-if)# **glbp 1 ip 192.168.1.1**



GLBP Configuration

Configure R1's Ethernet 0/1 with GLBP priority of 110 and enable preemption for both GLBP routers:

- R1(config)# **interface ethernet 0/1**
- R1(config-if)# **glbp 1 priority 110**
- R1(config-if)# **glbp 1 preempt**

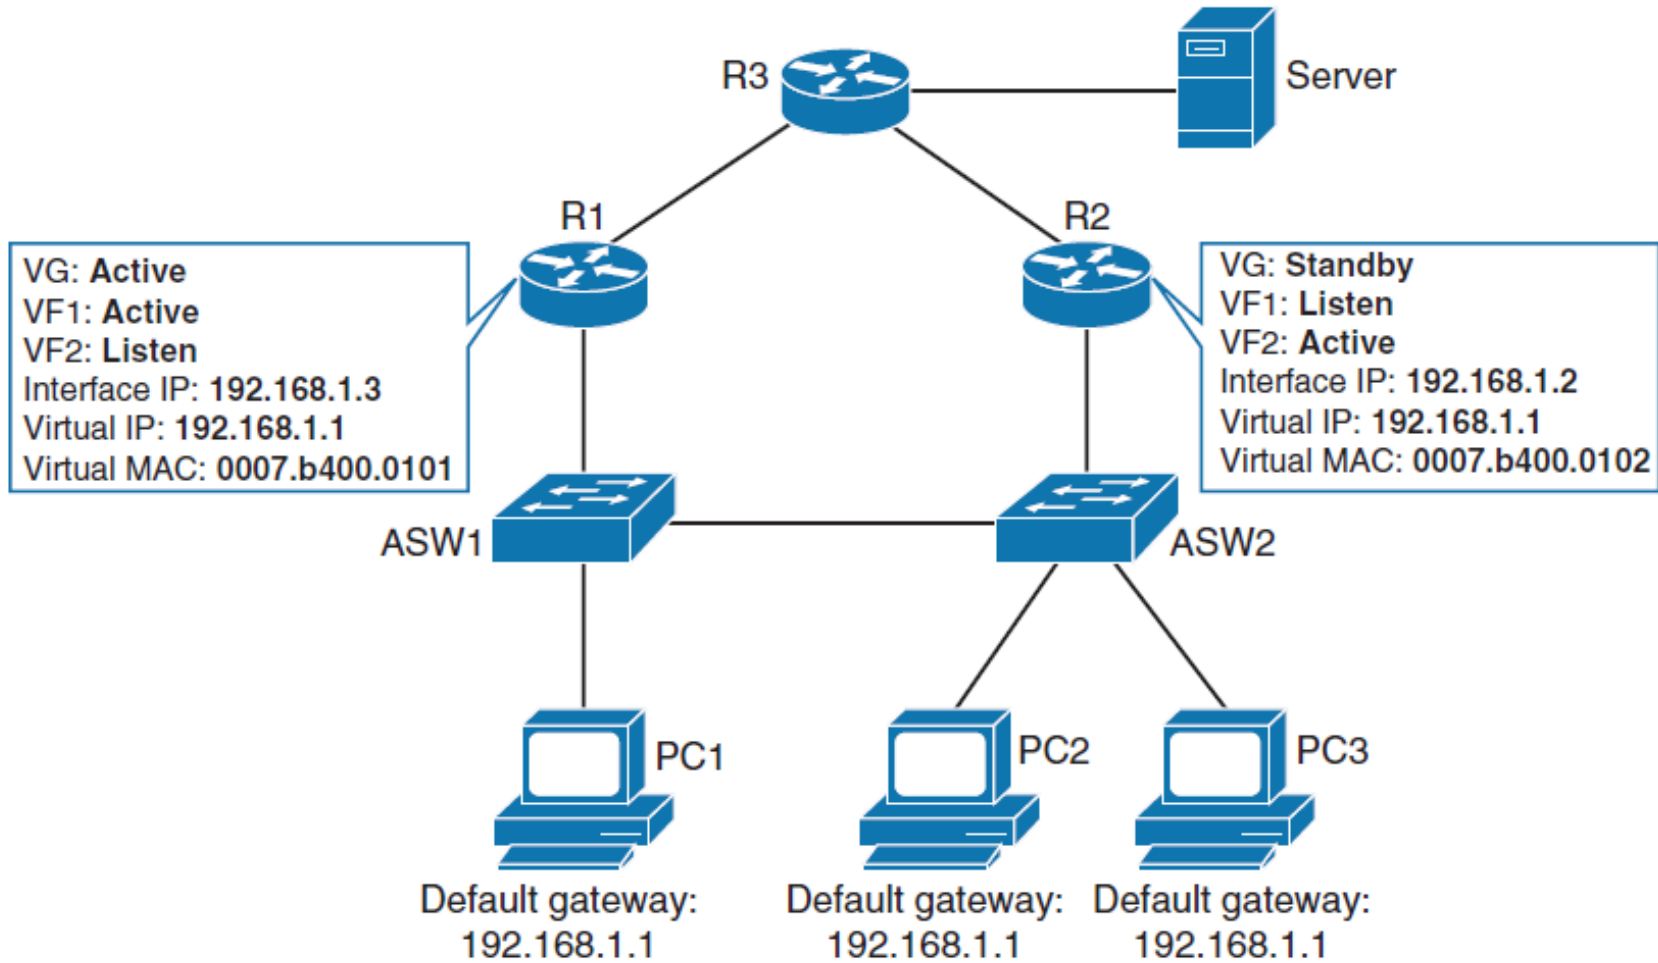
- R2(config)# **interface ethernet 0/1**
- R2(config-if)# **glbp 1 preempt**



The virtual MAC addresses of GLBP

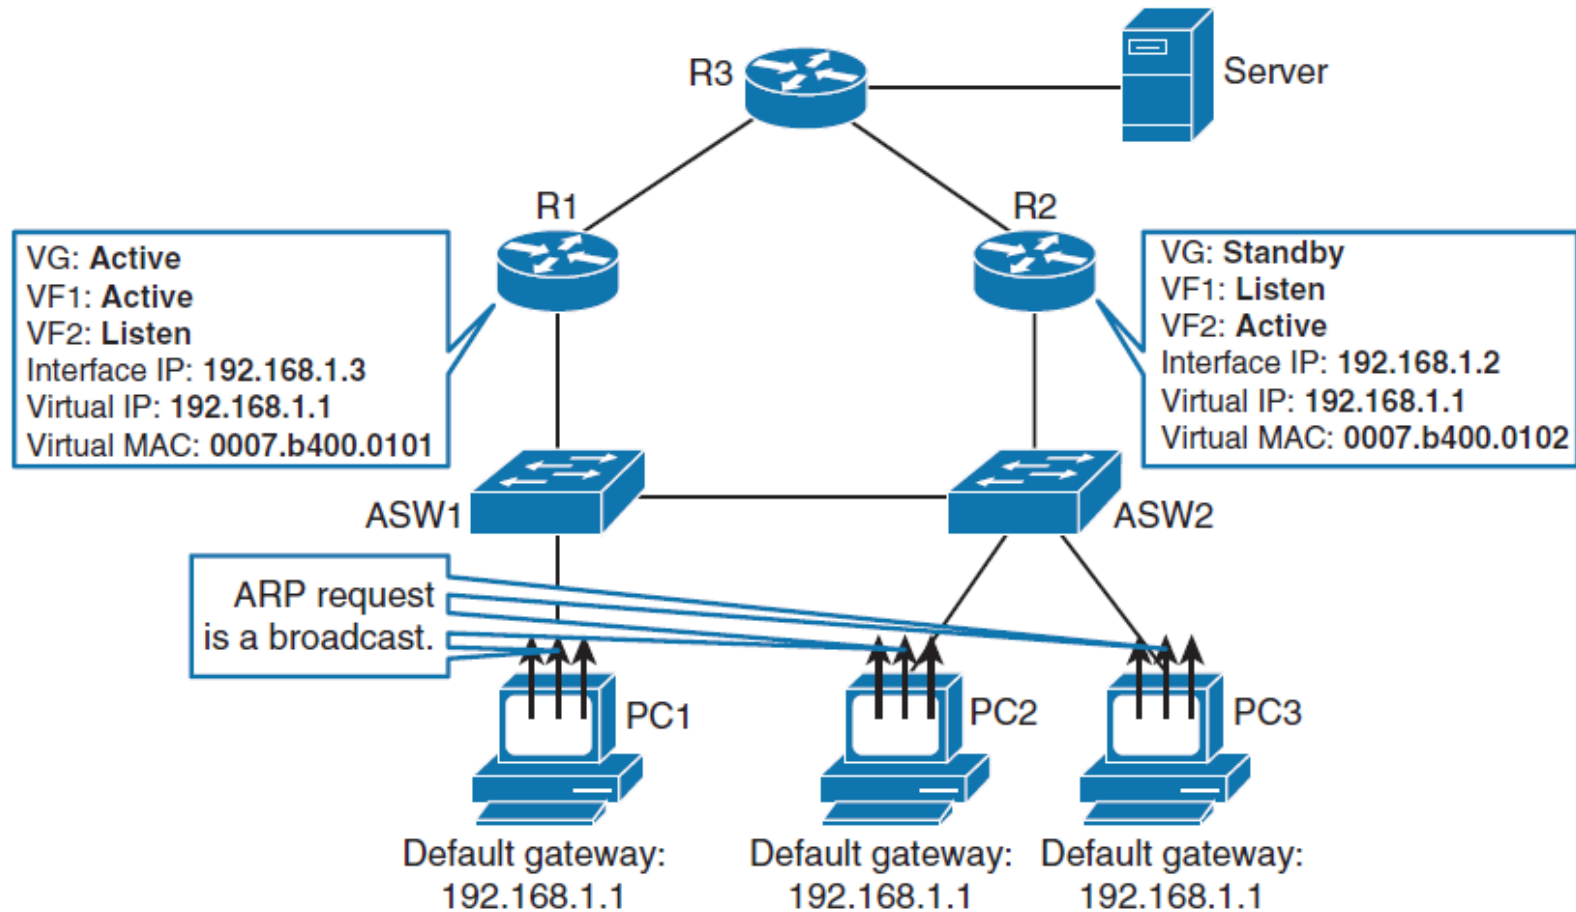
- The virtual MAC addresses of GLBP are in the form of 0007.b4XX.XXYY.
- XXXX is a 16-bit value that represents:
 - six 0 bits (6 x 0 bits),
 - followed by a 10-bit GLBP group number
- YY is an 8-bit value, and it represents the virtual forwarder number.
- The AVG assigned forwarder 1 virtual MAC address of 0007. b400.0101 and forwarder 2 virtual MAC address of 0007.b400.0102

GLBP Final Configuration



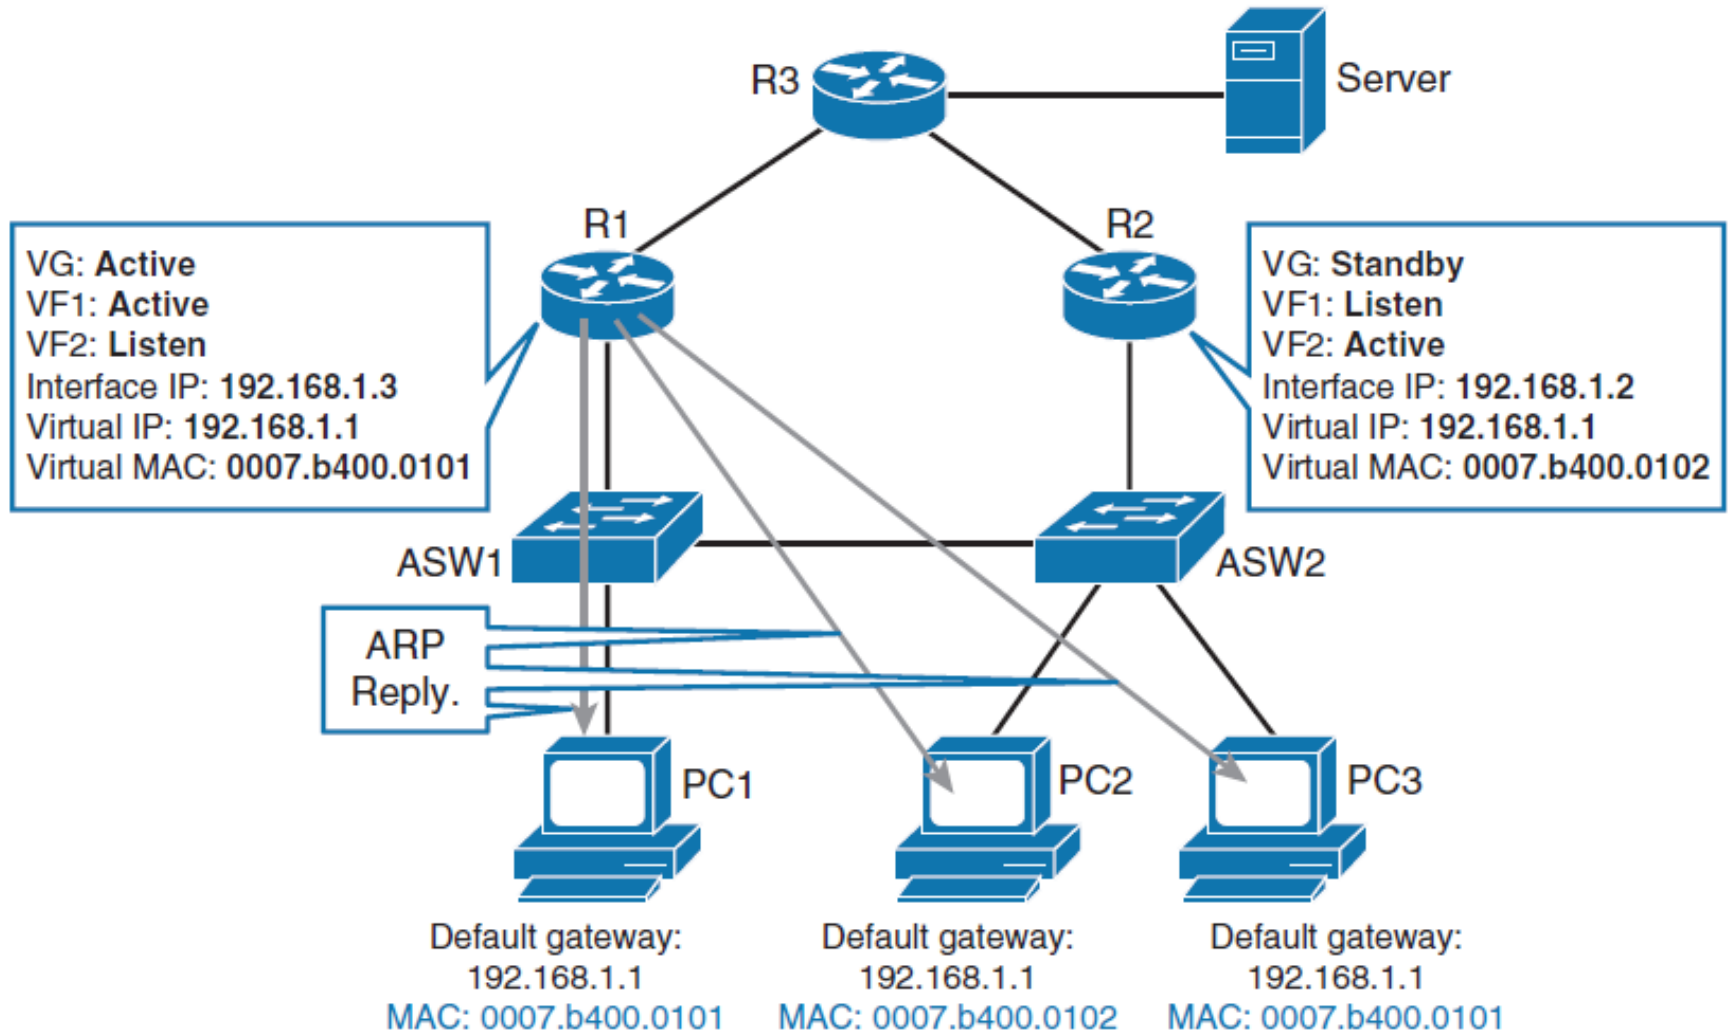


GLBP Operation (ARP Request)



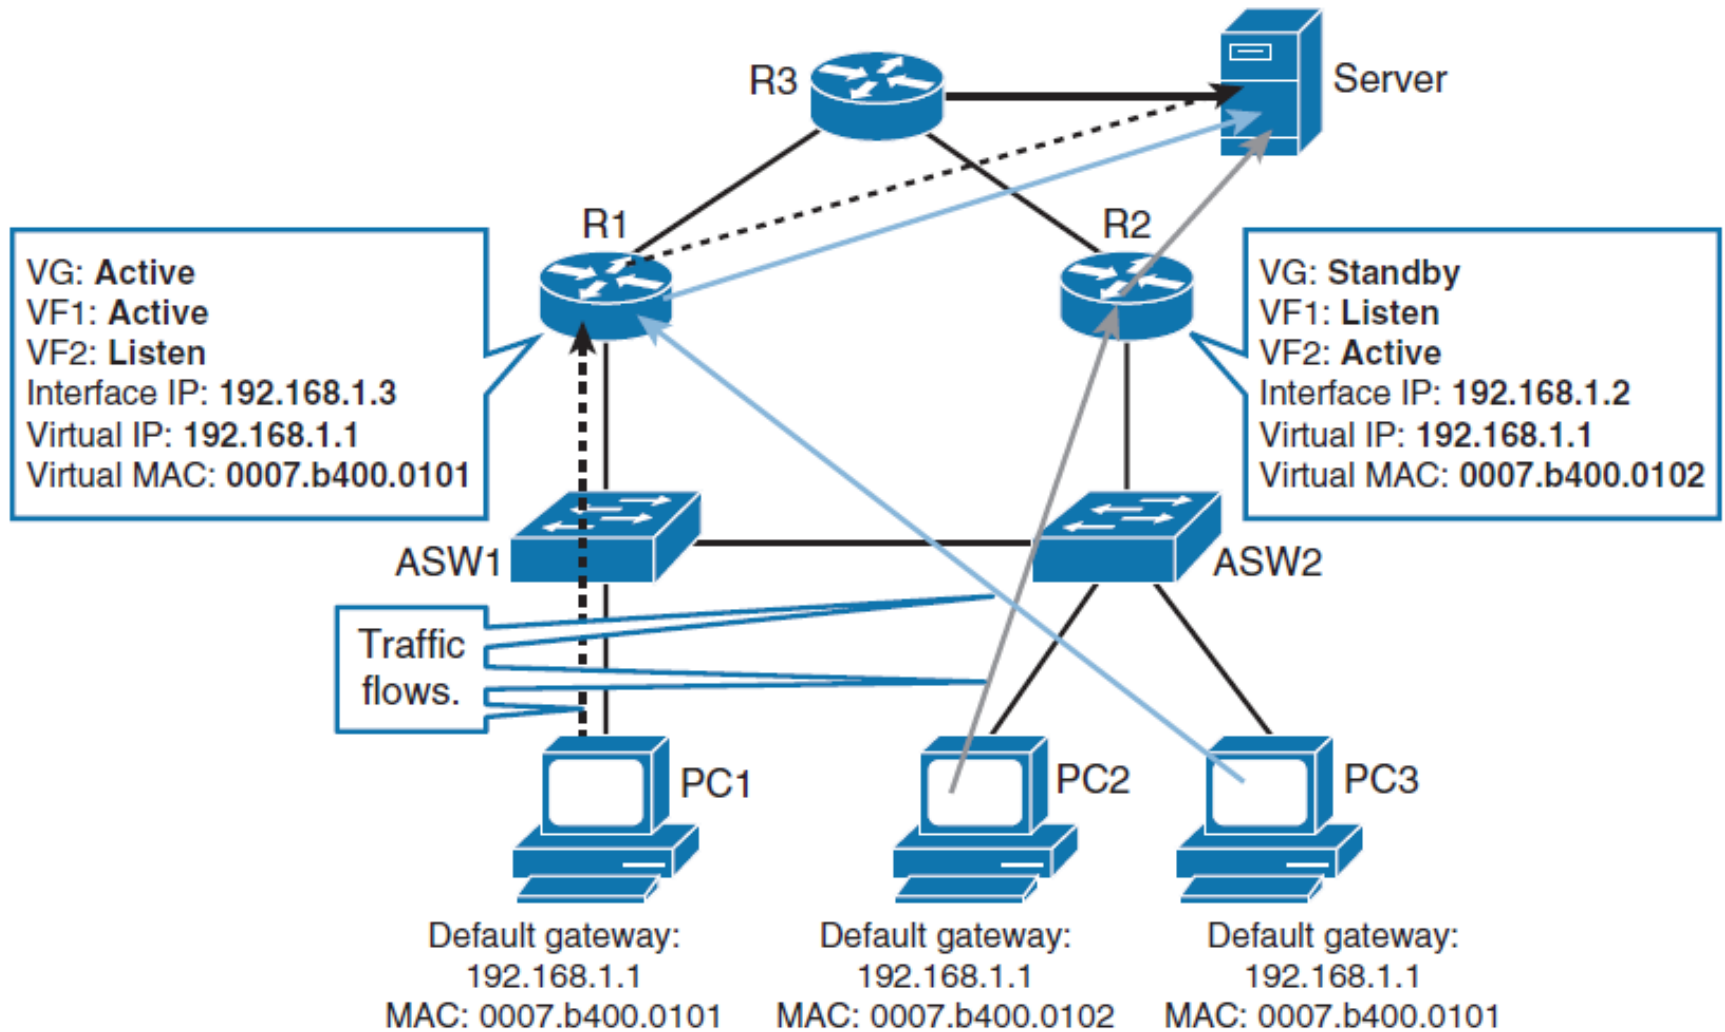


GLBP Operation (ARP Reply)



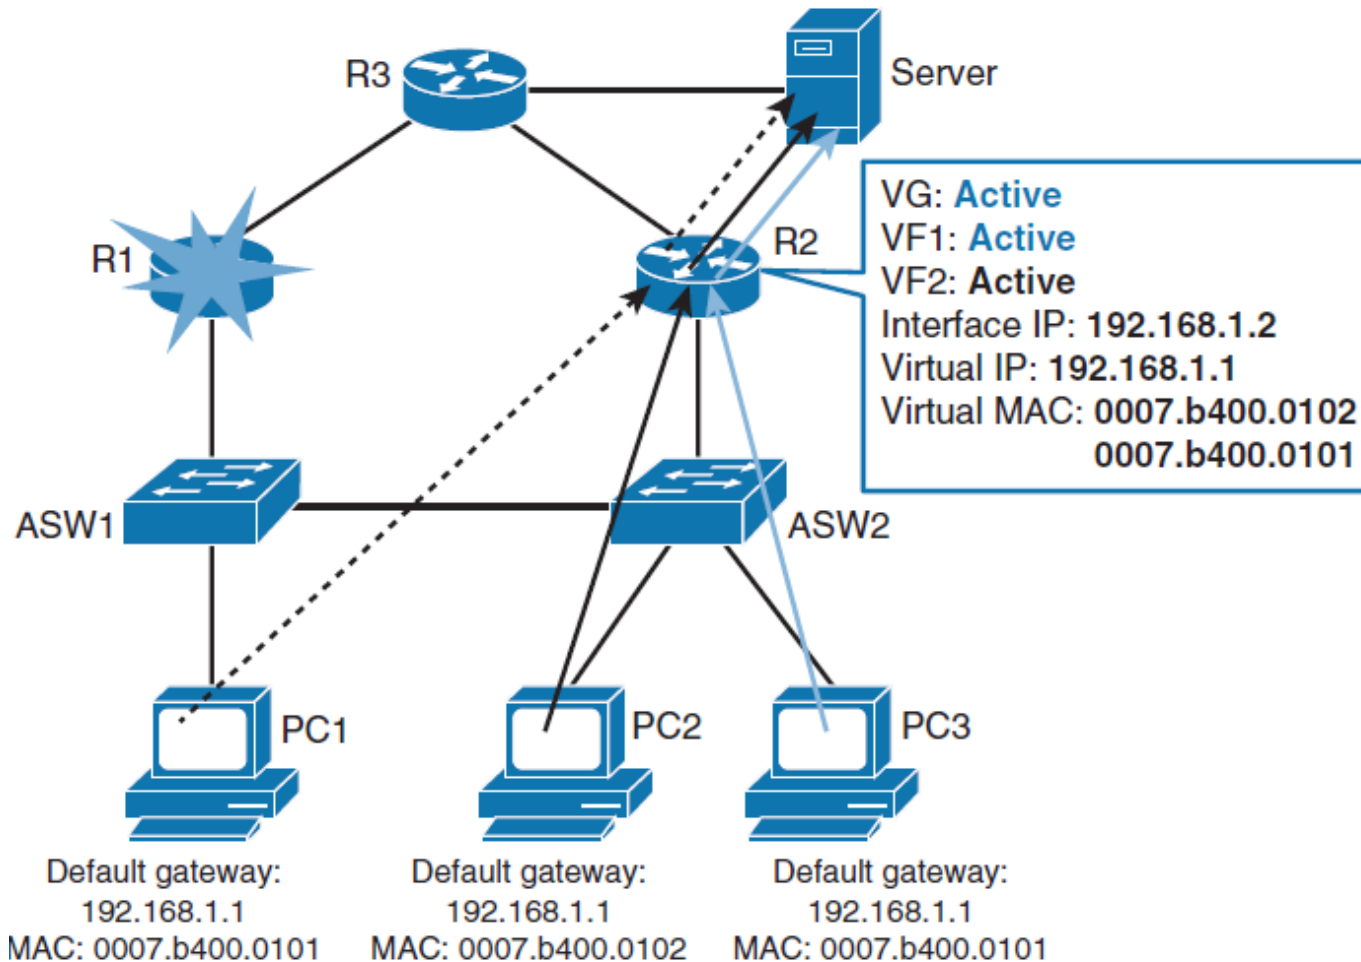


GLBP Operation (Traffic Flow)





GLBP Operations: Failed R1 New Data Path





GLBP Load-Balancing Options

GLBP supports the following operational modes for load balancing traffic across multiple default routers that are servicing the same default gateway IP address:

■ **Weighted load-balancing algorithm**

- The amount of load that is directed to a router depends on the weighting value that is advertised by that router.

■ **Host-dependent load-balancing algorithm**

- A host is guaranteed the use of the same virtual MAC address as long as that virtual MAC address is participating in the GLBP group.

■ **Round-robin load-balancing algorithm**

- As clients send ARP requests to resolve the MAC address of the default gateway, the reply to each client contains the MAC address of the next possible router in a round-robin fashion. The MAC addresses of all routers take turns being included in address resolution replies for the default gateway IP address.

■ To configure the load-balancing option, use the following command:

- `Switch(config-if)# glbp group load-balancing [round-robin | weighted | host-dependent]`



GLBP Authentication

- The key for the MD5 hash can either be given directly in the configuration using a key string or supplied indirectly through a key chain.
- The key string cannot exceed 100 characters in length.

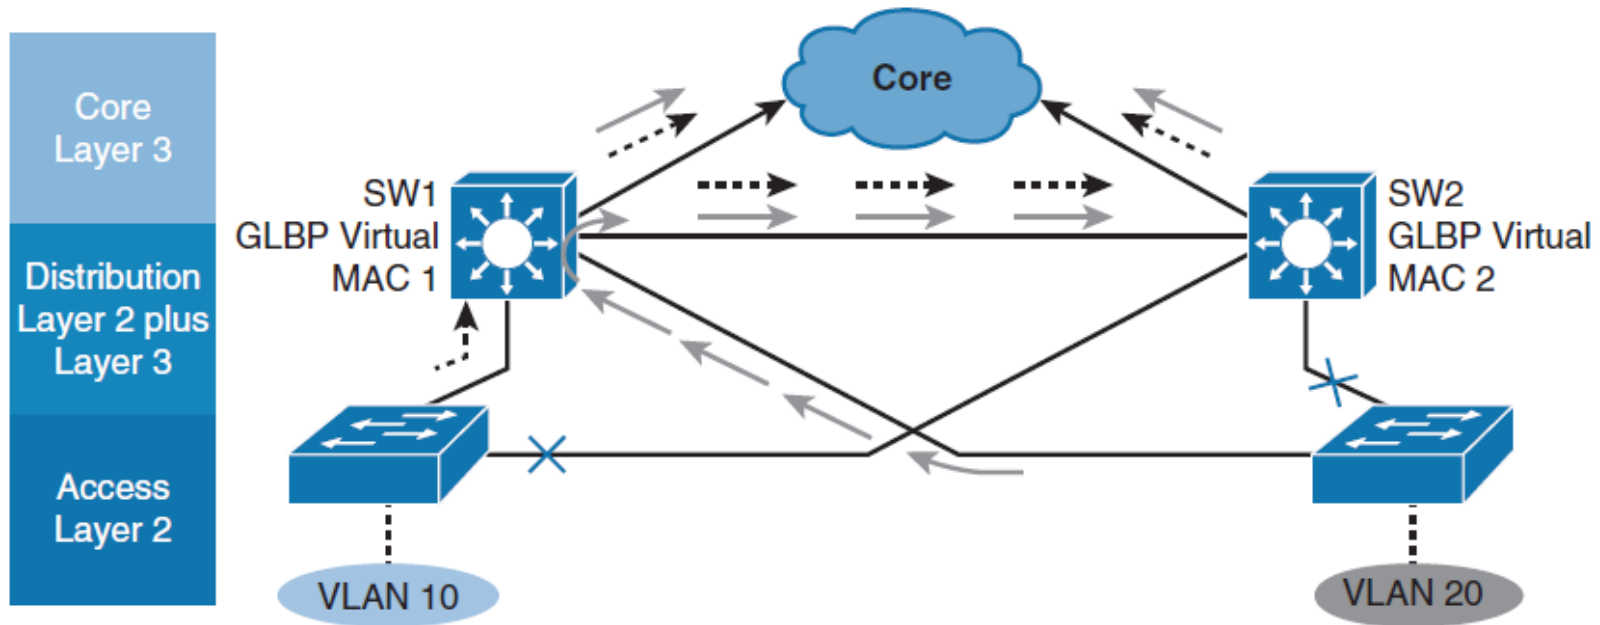
The following example demonstrates the configuration for GLBP authentication:

- Router(config)# **interface Ethernet0/1**
- Router(config-if)# **ip address 10.0.0.1 255.255.255.0**
- Router(config-if)# **glbp 1 authentication md5 key-string d00b4r987654321a**
- Router(config-if)# **glbp 1 ip 10.0.0.10**



GLBP and STP

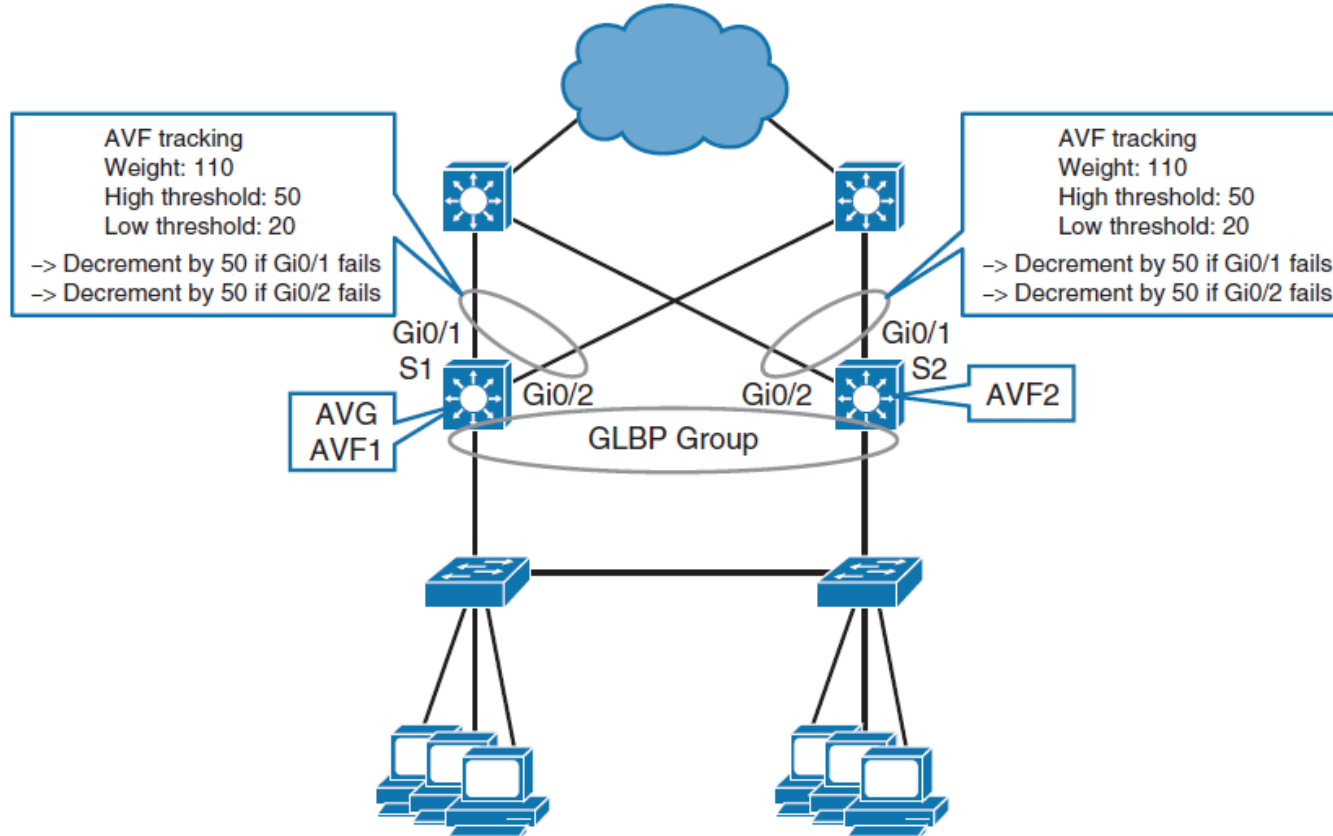
- With some switching topologies, the operation of STP results in inefficient traffic paths.
- In such cases, implementation of HSRP might be preferred over GLBP because it is easier to understand, whereas GLBP provides no advantages.





Tracking and GLBP

- Changing weight affects the AVF election and the load-balancing algorithm.
- Both values can be manipulated with object tracking.

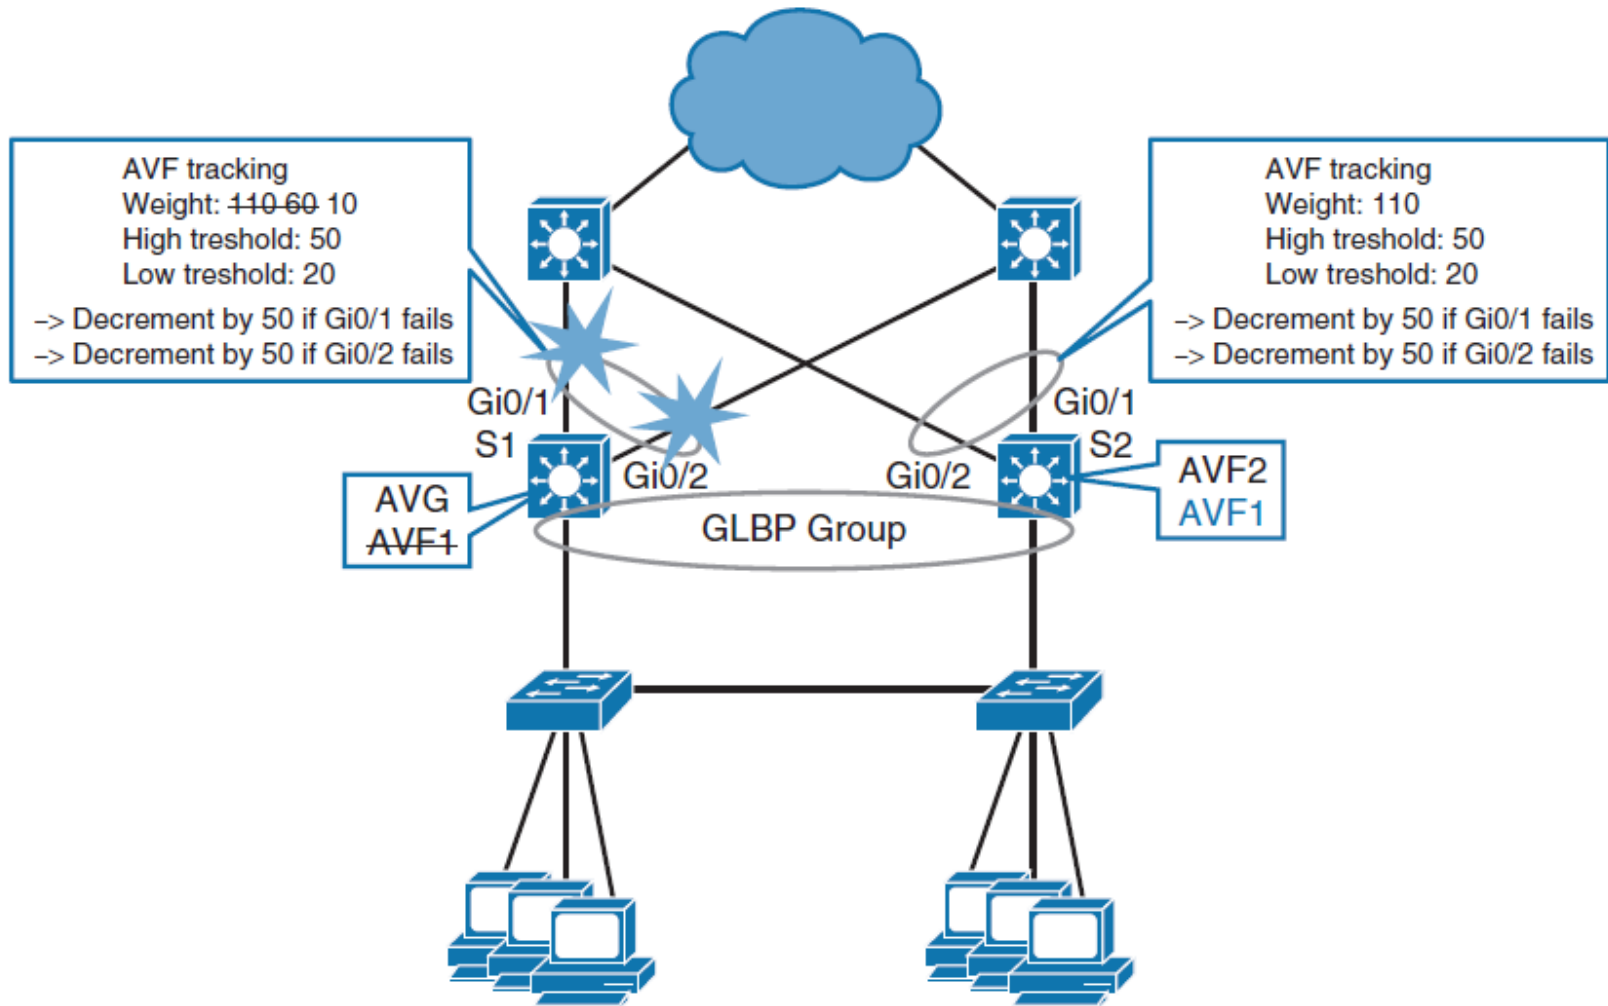




GLBP Weight

- GLBP uses a weighting scheme to determine the forwarding capacity of each router in the GLBP group.
- The weighting that is assigned to a router in the GLBP group can be used to determine whether it will forward packets and, if so, the proportion of hosts in the LAN for which it will forward packets.
- Thresholds can be set to disable forwarding when the weighting for a GLBP group falls below a certain value, and when it rises above another threshold, forwarding is automatically reenabled.
- By default, the GLBP virtual forwarder preemptive scheme is enabled with a delay of 30 seconds.
- A backup virtual forwarder can become the AVF if the current AVF weighting falls below the low weighting threshold for 30 seconds.
- To disable the GLBP forwarder preemptive scheme, use the **no glbp forwarder preempt** command or change the delay by using the **glbp forwarder preempt delay minimum** command.

GLBP Tracking Detects Interface Failure



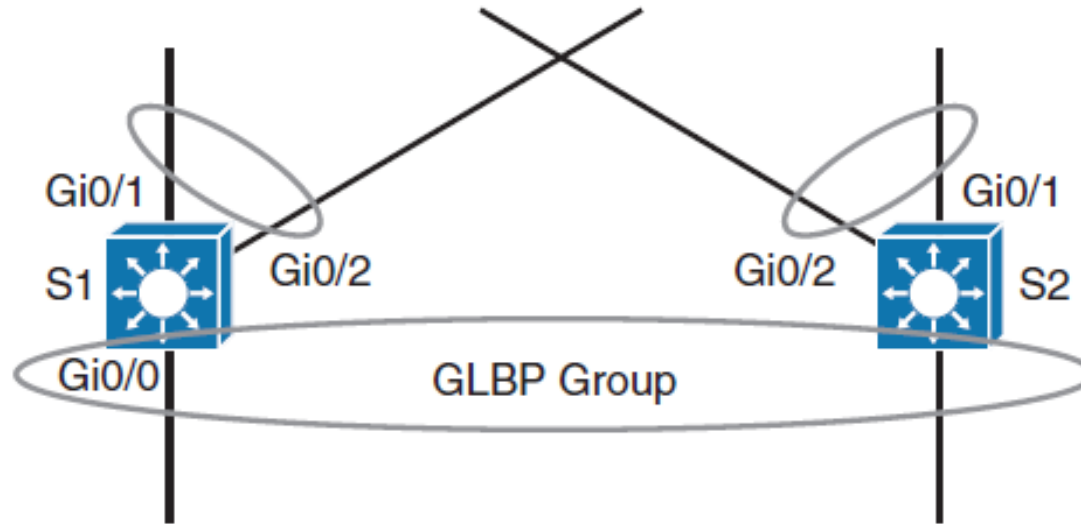


GLBP Weighing Option Under Failures

| Number of Active Uplinks on S1 | Number of Active Uplinks on S2 | Traffic Through S1 |
|--------------------------------|--------------------------------|-----------------------------------|
| 2 | 2 | 50% (110:110) |
| 2 | 1 | 65% (110:60) |
| 2 | 0 | 100% (S2 not eligible to forward) |
| 0 | 2 | 0% (S1 not eligible to forward) |
| 1 | 2 | 35% (60:110) |
| 1 | 1 | 50% (110:110) |
| 1 | 0 | 100% (S2 not eligible to forward) |
| 0 | 1 | 0% (S1 not eligible to forward) |
| 0 | 0 | 0% (none eligible for forwarding) |



GLBP Object Tracking Sample Configuration



```

s1(config)# track 1 interface GigabitEthernet0/1 line-protocol
s1(config)# track 2 interface GigabitEthernet0/2 line-protocol
s1(config-track)# interface GigabitEthernet0/0
s1(config-if)# ip address 192.168.1.2 255.255.255.0
s1(config-if)# glbp 1 ip 192.168.1.1
s1(config-if)# glbp 1 priority 110
s1(config-if)# glbp 1 preempt
s1(config-if)# glbp 1 weighting 110 lower 20 upper 50
s1(config-if)# glbp 1 weighting track 1 decrement 50
s1(config-if)# glbp 1 weighting track 2 decrement 50
    
```



Chapter 6 Summary

- The redundancy protocol provides the mechanism for determining which router should take the active role in forwarding traffic and determining when that role must be taken over by a standby router .
- HSRP is a Cisco proprietary protocol, whereas VRRP is an industry standard for virtual routing gateways.
- HSRP Version 1 and Version 2 active and standby routers send hello messages to multicast address 224.0.0.2 for Version 1 and 224.0.0.102 for Version 2 on UDP port 1985.
- It is important that the configured active router should be the same as the STP root bridge.
- HSRP and VRRP use the VLAN load-balancing mechanism for load balancing.
- With the new RFC, only the Cisco implementation of VRRP supports VRRP authentication.
- GLBP, by default, provides the virtual gateway and load balancing via multiple virtual MAC addresses.
- Review all the configuration examples and troubleshooting steps for better understanding and for exam preparation.



Chapter 6 Labs

- **CCNPv7.1 SWITCH Lab6.1 FHRP HSRP VRRP**
- **CCNPv7.1 SWITCH Lab6.2 HSRPv6**
- **CCNPv7.1 SWITCH Lab6.3 GLBP**

Cisco | Networking Academy[®]

Mind Wide Open[™]



Acknowledgment

- *Some of the images and texts are from Implementing Cisco IP Switched Networks (SWITCH) Foundation Learning Guide: (CCNP SWITCH 300-115) by Richard Froom and Erum Frahim (1587206641)*
- Copyright © 2015 – 2016 Cisco Systems, Inc.
- Special Thanks to *Bruno Silva*