

## 3 Approaches to router redundancy

HSRP – Hot Standby Router Protocol

VRRP – Virtual Router Redundancy Protocol

GLBP – Gateway Load Balancing Protocol

*Why is redundancy important?* Without an available gateway, packets cannot be forwarded from their VLAN or subnet.

### Hot Standby Router Protocol

HSRP is a Cisco-proprietary protocol. More information can be found in RFC2281.

HSRP allows 2 or more routers (or multilayer switches) to appear as a single gateway. Routers are assigned to a HSRP group. 1 router is elected the primary or **active** HSRP router. A second router is elected the **standby** HSRP router. The remaining routers are placed into the **listen** state. All the routers in a HSRP group exchange **hello** messages at regular intervals.

HSRP **hello** messages are sent to 224.0.0.2 (“all routers”) on UDP port 1985

HSRP groups are assigned a number from 0-255. Most Catalyst switches only support a maximum of 16 unique HSRP groups. However these groups are only significant on a per interface basis. You can use group 1 each interface you are configuring for HSRP and only use 1 unique HSRP group total.

HSRP groups are configured under the interface you would like to use HSRP on

### HSRP Router Election

Routers are elected based on a priority value from 0-255, which is configured on each router in the HSRP group. The highest priority (255) wins becoming the active router for the group. If the priorities are all equal, the router with the highest IP address wins. The default priority is 100.

Command to set the priority:

```
switch (config-if)# standby group priority priority  
switch (config-if)# standby 1 priority 150
```

When HSRP is configured on an interface, the router goes through a series of states before it reaches the **active** state. The HSRP states order is:

1. Disabled
2. Init
3. Listen
4. Speak
5. Standby
6. Active

Only the **standby** router monitors hello messages from the **active** router. Hello messages are sent every 3 seconds by default. If hello messages are missed for the duration of the **hold-down timer** (the default is 10 seconds or 3x the hello timer), the **active** router is presumed to be down and the **standby** router can assume the **active** role. If there are routers in the **listen** state, the router with the next highest priority becomes the **standby** router for the HSRP group.

The timers can be adjusted manually by using the following command:

```
switch (config-if)# standby group timers [msec] hello [msec] hold-down  
switch (config-if)# standby 1 timers 1 4  
switch (config-if)# standby 1 timers msec 100 msec 400
```

The timer values can be given in either seconds or milliseconds.

Normally after the **active** router fails and recovers, it cannot become the **active** router again, even if it has the highest priority. For this to happen, you must use enable preemption.

```
switch (config-if)# standby group preempt [delay [minimum secs] [reload secs]]  
switch (config-if)# standby 1 preempt  
switch (config-if)# standby 1 preempt delay minimum 60  
switch (config-if)# standby 1 preempt delay reload 120
```

The **minimum** command forces the router to wait between 0 and 3600 seconds before assuming the **active** role.

The **reload** option forces the router to wait between 0 and 3600 seconds before assuming the **active** role after it has been reloaded or restarted. This handy if you need it to wait for routing protocols to converge.

HSRP can also use authentication to prevent unexpected devices from participating in the router election. If used, it must be configured on all interfaces in the HSRP group.

### Plain-text HSRP Authentication

This form of authentication uses a plain-text key string up to 8 characters long to authenticate HSRP group peers. Sends the key as clear text and should only be used to prevent HSRP group peers with a default configuration from participating.

```
switch (config-if)# standby group authentication string  
switch (config-if)# standby 1 authentication alphabet
```

### MD5 Authentication

An MD5 hash is computed on a portion of each HSRP message and a secret key known only to the HSRP group peers. This hash is sent with each HSRP message. As messages are received, the receiving peer recomputes the hash of the message and key. If they match, the peer accepts the message.

MD5 authentication is more secure than plain-text because only the hash is sent and not the key.

You configure MD5 authentication by associating a key string with the HSRP interface:

```
switch (config-if)# standby group authentication md5 key-string [0 | 7] string
switch (config-if)# standby 1 authentication md5 key-string supersecretkey
```

The key string can be up to 64 characters in length and is given as plain-text by default. This is the same as using option 0. You can also cut and paste a value by using option 7.

You can define an MD5 key string as a key on a key chain. This is much more flexible, allowing you to define more than one key on the switch. If you need to change a key you can just add a new key and then delete the old one when you are ready.

To configure MD5 authentication using a key chain:

```
switch (config)# key chain chain-name
switch (config-keychain)# key key-number
switch (config-keychain-key)# key-string [0 | 7] string
switch (config)# interface type mod/num
switch (config-if)# standby group authentication md5 key-chain chain-name

switch (config)# key chain mykeychain
switch (config-keychain)# key 100
switch (config-keychain-key)# key-string supersecretkey
switch (config)# interface fa0/1
switch (config-if)# standby 1 authentication md5 key-chain mykeychain
```

## Conceding the Election

Because a gateway is required to reach the rest of the world, an **active** HSRP router must be able to concede the **active** role when it is no longer able to reach other networks.

HSRP can be configured to track interfaces and reduce a router's priority by a configured amount when the tracked interface goes down. It also allows the priority to be increased once the interface is available again. This happens immediately.

To configure tracking on an interface:

```
switch (config-if)# standby group track type mod/num [decrementvalue]
switch (config-if)# standby 1 track fa0/1 25
```

By default the *decrementvalue* is 10. Remember, interface tracking involves the state useful router interfaces and not the state of the HSRP interface. The **standby** router can only take over and become the **active** router if the following conditions are met:

- it has a higher HSRP priority, **and**
- it has **preempt** enabled in the HSRP configuration

As in any other router election, if a router has a higher HSRP priority than the current **active** router and preemption is not enabled, the **standby** router will not assume the **active** router position.

## HSRP Gateway Addressing

Each router interface has unique IP address used for routing protocols and management traffic directed to or initiated by the router. Each router also shares a common gateway IP address, the virtual address, which is used by HSRP. This is known as the HSRP address or standby address. This is the address clients should use as their default gateway to always have one available. The actual and virtual IP addresses must be in the same subnet.

To configure an HSRP address:

```
switch (config-if)# standby group ip ip_addr [secondary]
switch (config-if)# standby 1 ip 192.168.1.1
```

The secondary option is used to provide a redundant secondary gateway address on an interface configured with a secondary address.

The MAC address for the HSRP address is in the form of 0000.0c07.acXX where XX is the HSRP group number.

```
Group 1 is 0000.0c07.ac01
Group 2 is 0000.0c07.ac02
...
Group 16 is 0000.0c07.ac10
```

### Sample HSRP configuration and explanation

```
01 switch1> en
02 switch1# conf t
03 switch1 (config)# interface vlan 10
04 switch1 (config-if)# description *** User Network ***
05 switch1 (config-if)# ip address 192.168.10.10 255.255.255.0
06 switch1 (config-if)# standby 1 priority 150
07 switch1 (config-if)# standby 1 preempt
08 switch1 (config-if)# standby 1 ip 192.168.10.1

09 switch2> en
10 switch2# conf t
11 switch2 (config)# interface vlan 10
12 switch2 (config-if)# description *** User Network ***
13 switch2 (config-if)# ip address 192.168.10.11 255.255.255.0
14 switch2 (config-if)# standby 1 priority 125
15 switch2 (config-if)# standby 1 preempt
16 switch2 (config-if)# standby 1 ip 192.168.10.1
```

## Load Balancing with HSRP

Load balancing with HSRP requires that 2 groups be configured on each router. 1 group is configured as the **active** router for the first group and the second router is configured as the **active** router for the second group. The clients are then configured so they half use one HSRP address as their gateway and

the remaining clients use the second HSRP address as their gateway. Then should a router fail, the gateway is still available.

### Sample configuration for HSRP load balancing

```
switch1> en
switch1# conf t
switch1 (config)# interface vlan 10
switch1 (config-if)# description *** User Network ***
switch1 (config-if)# ip address 192.168.10.10 255.255.255.0
switch1 (config-if)# standby 1 priority 200
switch1 (config-if)# standby 1 preempt
switch1 (config-if)# standby 1 ip 192.168.10.1
switch1 (config-if)# standby 2 priority 125
switch1 (config-if)# standby 2 preempt
switch1 (config-if)# standby 2 192.168.10.2

switch2> en
switch2# conf t
switch2 (config)# interface vlan 10
switch2 (config-if)# description *** User Network ***
switch2 (config-if)# ip address 192.168.10.11 255.255.255.0
switch2 (config-if)# standby 1 priority 125
switch2 (config-if)# standby 1 preempt
switch2 (config-if)# standby 1 ip 192.168.10.1
switch2 (config-if)# standby 2 priority 200
switch2 (config-if)# standby 2 preempt
switch2 (config-if)# standby 2 192.168.10.2
```

With the above configuration, you would point half the network clients to the 192.168.10.1 address for their gateway and the remaining would use 192.168.10.2 as their gateway.

To show the status of your configured HSRP group(s) you use the following command

```
switch1# show standby [brief] [vlan vlan-id | type mod/num]

switch1# show standby brief vlan 10

switch1# show standby vlan 10

switch1# show standby brief fa0/1

switch1# show standby fa0/1
```

The information provided by the **brief** option includes the which router is the active and which is the standby, the actual IP addresses and the HSRP addresses.

The show command without the **brief** option displays the above information and other information including the HSRP timers and authentication text. This command is the basic command used to verify your HSRP groups are configured correctly and the appropriate router is active when troubleshooting network issues.

## Virtual Router Redundancy Protocol

VRRP is the standards based alternative to HSRP and defined in the IETF standard RFC2338.

Differences in VRRP from HSRP

- Provides 1 redundant address from a group. **Active** router is called the **master**, and the rest are in the **backup** state. **Master** router has the highest priority in the group
- VRRP group ranges 0-255
- VRRP priorities range 1-254 with 100 being the default and 254 being the highest
- Virtual router MAC address is 0000.5e00.01xx where xx is the group number
- VRRP advertisements are sent every 1 second. Backup routers can learn timer settings from the **master** router
- **backup** routers can preempt the **master** router by default
- No way to track interfaces using VRRP

The essential VRRP commands

```
switch(config-if)# vrrp group priority priority
switch(config-if)# vrrp group timers advertise [msec] interval
switch(config-if)# vrrp group timers learn
switch(config-if)# no vrrp group preempt
switch(config-if)# vrrp group preempt [delay seconds]
switch(config-if)# vrrp group authentication string
switch(config-if)# vrrp group ip ip_address [secondary]
switch# show vrrp [brief]
```

Sample load balancing configuration using VRRP

```
switch1(config)# int vlan 10
switch1(config-if)# ip add 192.168.10.10 255.255.255.0
switch1(config-if)# vrrp 1 priority 200
switch1(config-if)# vrrp 1 ip 192.168.10.1
switch1(config-if)# vrrp 2 priority 100
switch1(config-if)# no vrrp 2 preempt
switch1(config-if)# vrrp 2 ip 192.168.10.2

switch2(config)# int vlan 10
switch2(config-if)# ip add 192.168.10.11 255.255.255.0
switch2(config-if)# vrrp 1 priority 100
switch2(config-if)# no vrrp 1 preempt
switch2(config-if)# vrrp 1 ip 192.168.10.1
switch2(config-if)# vrrp 2 priority 200
```

```
switch2(config-if)# vrrp 2 ip 192.168.10.2
```

## Gateway Load Balancing Protocol

GLBP is a Cisco-proprietary protocol created to overcome the limitations of existing redundancy protocols. The concepts are similar, but behavior of GLBP is much more dynamic.

To create a virtual router, multiple routers (or MLS) are assigned to a GLBP group. Rather than a single active router, each router in the group can participate, forwarding traffic and offer load balancing. The advantage of GLBP is that all clients can be pointed to the same default gateway reducing administration. Load balancing is provided through the use of virtual router MAC addresses.

### Active Virtual Gateway

1 router is elected the active virtual gateway (AVG). The AVG has the highest priority or highest IP address if the priorities match. The AVG answers all ARP requests and returns the MAC address based on the load balancing algorithm that is configured.

The active virtual gateway can assign up to 4 virtual MAC addresses in a group. Each is referred to as an active virtual forwarder (AVF). Other routers in the group serve as backup virtual forwarders. These roles are also assigned by the AVG.

To assign a GLBP interfaces priorities

```
switch(config-if)# glbp group priority priority  
switch(config-if)# glbp 1 priority 200
```

GLBP groups range from 0-1023 and the priorities can range 1-255 with 100 being the default value. Preemption works the same way in GLBP as it does when using HSRP, the current active router must fail before the standby router will take over the role.

To enable preemption on an interface

```
switch(config-if)# glbp group preempt [delay seconds]  
switch(config-if)# glbp 1 preempt delay 10
```

### Active Virtual Forwarder

Each router in a GLBP groups sends hellos to every other GLBP router. If an AVF fails, the AVG assigns the virtual router MAC to another AVF in the group. The AVG then waits for the duration of the **redirect** timer to expire. Once this has expired the AVG no longer responds to ARP requests for the failed AVF. After the duration of the **timeout** timer has passed, it is assumed the failed AVF is no longer expected to come back online.

To configure the timers

```
switch(config-if)# glbp group timers redirect redirect timeout  
switch(config-if)# glbp 1 timers redirect 300 3660
```

The redirect default is 600 seconds and the timeout default is 4 hours.

Like HSRP, GLBP can track interfaces and assign a weight to decrement the priority by so less useful routers can drop from the GLBP group.

```
switch(config)# track object-number interface type mod/num {line-protocol | ip
routing}
switch(config-if)# glbp group weighting maximum [lower lower] [upper upper]
switch(config-if)# glbp group weighting track object-number [decrement value]

switch(config)# track 10 interface fa0/2 line-protocol
switch(config-if)# glbp 1 weighting 120 lower 100
switch(config-if)# glbp 1 weighting track 10 decrement 25
```

add explanation here.

## GLBP Load Balancing

The AVG load balances traffic by handing out virtual router MAC addresses in a deterministic fashion.

Load balancing methods available in a GLBP group are:

- Round robin – this is the default
- Weighted – based on interface weight. Interfaces with higher weight receive more ARP requests
- Host-dependent – each client ARP request always receives the same MAC address

To configure load-balancing on the AVG (or successors)

```
switch(config-if)# glbp group load-balancing [round-robin | weighted | host-
dependent]
switch(config-if)# glbp 1 load-balancing host-dependent
```

To enable GLBP you need to assign a virtual router IP address to the group

```
switch(config-if)# glbp group ip [ip_address [secondary]]
switch(config-if)# glbp 1 ip 192.168.10.1
```

## Redundancy Within a Switch Chassis

Some Cisco switches can provide redundancy for the supervisor as long as redundant hardware is in place. This is also possible with the power supply.

### Redundant Switch Supervisors

Catalyst 4500R and 6500 can accept 2 supervisors. The first supervisor to boot becomes the active supervisor. All switching functions are provided by the active supervisor.

The available modes of redundancy are

- Route Processor Redundancy (RPR) – redundant supervisor is only partially booted and

initialized. Must reload every other module in the switch when it becomes active

- Route Processor Redundancy Plus (RPR+) - the redundant supervisor is fully, however no layer 2 or 3 functions are running. It does allow the switchports to maintain their current state
- Stateful Switchover (SSO) – redundant supervisor is fully booted and initialized. Layer 2 information is maintained on both supervisors for continued switching during failover. Also allows for ports to maintain current state.

## Configuring Redundancy

```
router(config)# redundancy
router(config-red)# mode {rpr | rpr-plus | sso}
```

For rpr-plus to work, the IOS images of the supervisors must match.

## Configuring Supervisor Synchronization

By default the start-up configuration and configuration register are synchronized with the standby supervisor.

To configure other options:

```
router(config)# redundancy
router(config-red)# main cpu
router(config-r-c)# auto-sync {startup-config | config-register | bootvar}
```

You can use the auto-sync command multiple times to select more than one option. You can also use auto-sync standard to return to the defaults.

## Non-Stop Forwarding

This feature works with SSO and focuses on rebuilding the routing information base quickly. It is a Cisco-proprietary function and supports the BGP, EIGRP, OSPF and IS-IS routing protocols.

NSF must be configured on the router needing assistance and the router(s) providing assistance.

## Redundant Power Supplies

Available on the 6500 and 4500R platforms. The power supplies must be identical for redundancy to be configured and work.

Two modes are available:

- Combined mode – both power supplies work together and share the load
- Redundant mode – work independently, but only is actively supplying power at a time

To configure power redundancy

```
switch(config)# power redundancy-mode {redundant | combined}
switch(config)# power redundancy-mode combined
```

To control power to a module

```
switch(config)# [no] power enable module slot  
switch(config)# power enable module 8
```

Using the no option disables power to the selected module.