


- Redundancy and Load Balancing -

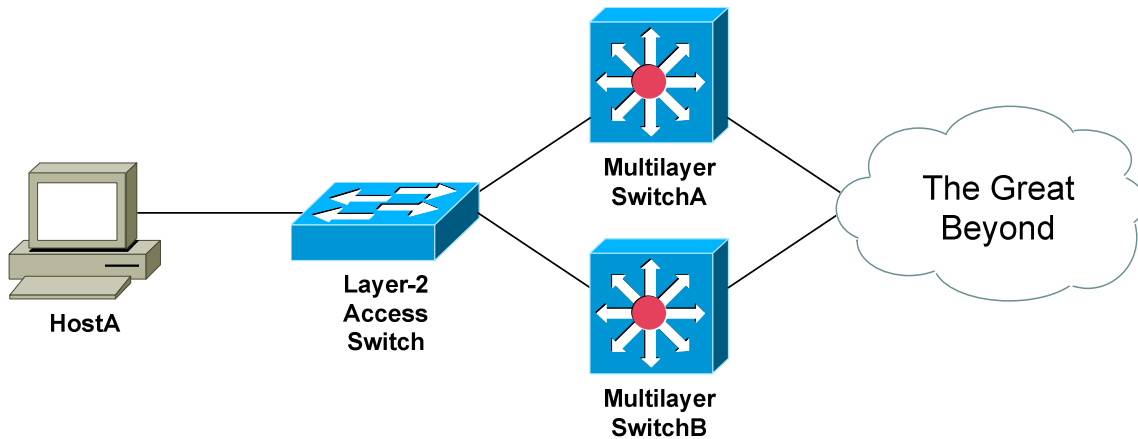
Importance of Redundancy

High availability is critical in most environments. Even a brief outage due to hardware failure may be considered unacceptable. Consider the following example:



To reach other networks, HostA must utilize a single *gateway* – SwitchA. The gateway represents a **single point of failure** on this network. If the gateway fails, hosts will lose access to *all* resources beyond the gateway.

Using multiple routers or multilayer switches can provide **Layer-3 redundancy** for hosts:



However, the Layer-3 redundancy must be **transparent** to each host. Hosts should not be configured with multiple default gateways.

Cisco supports three protocols to provide transparent Layer-3 redundancy:

- **Hot Standby Router Protocol (HSRP)**
- **Virtual Router Redundancy Protocol (VRRP)**
- **Gateway Load Balancing Protocol (GLBP)**

Note: The terms *multilayer switch* and *router* will be used interchangeably throughout this guide.

* * *

All original material copyright © 2014 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

Hot Standby Router Protocol (HSRP)

Cisco developed the proprietary **Hot Standby Router Protocol (HSRP)** to allow multiple routers or multilayer switches to masquerade as a single gateway. This is accomplished by assigning a **virtual IP and MAC address** to all routers participating in an HSRP **group**.

Routers within the same HSRP group must be assigned the same **group number**, which can range from **0 to 255**. However, most Cisco platforms only support **16** configured HSRP groups.


HSRP routers are elected to specific roles:

- **Active Router** – router currently serving as the gateway.
- **Standby Router** – backup router to the Active Router.
- **Listening Router** – all other routers participating in HSRP.

Only **one active** and **one standby** router are allowed per HSRP group. Thus, HSRP provides *Layer-3 redundancy*, but no inherent *load balancing*.

Hello packets are used to elect HSRP roles and to ensure all routers are functional. If the current active router fails, the standby router will immediately take over as active, and a new standby is elected. By default, hello packets are sent every **3 seconds**.

The role of an HSRP router is dictated by its **priority**. The priority can range from 0 – 255, with a default of **100**. A *higher* priority is preferred.



Thus, the router with the **highest priority** is elected the active router – *SwitchB* in the above example. The router with the **second highest priority** becomes the standby router – *SwitchA* in the example. If all priorities are equal, whichever router has the **highest IP Address** on its HSRP interface is elected the active router.

* * *

All original material copyright © 2014 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

HSRP States

A router interface participating in HSRP must progress through several **states** before settling into a role:

- **Disabled**
- **Initial**
- **Learn**
- **Listen**
- **Speak**
- **Standby**
- **Active**

A **disabled state** indicates that the interface is either not configured for HSRP, or is administratively shutdown.

An interface begins in an **initial state** when first configured with HSRP, or taken out of an administratively shutdown state.

An interface enters a **learn state** if it does not know the HSRP virtual IP address. Normally the virtual IP is manually configured on the interface – otherwise, it will be learned from the current Active router via hello packets.

An interface in a **listen state** knows the virtual IP address, but was not elected as either the Active or Standby Router.

Interfaces in a **speak state** are currently participating in the election of an active or standby router. Elections are performed using hello packets, which are sent out every 3 seconds by default.

A **standby state** indicates that the interface is acting as a backup to the active router. The standby router continuously exchanges hello packets with the active router, and will take over if the active router fails.

An interface in an **active state** is the live gateway, and will forward traffic sent to the virtual IP address. Hosts will use the virtual IP address as their **default gateway**. The active router will respond to ARP requests for the virtual IP with the virtual MAC address.

Note that hello packets are only exchanged in three HSRP states:

- **Speak**
- **Standby**
- **Active**

Interfaces in a **listen state** will only *listen* for hello packets. If an active or standby router fails, a listen interface will transition to a speak state to participate in a new election.

* * *

All original material copyright © 2014 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

HSRP Basic Configuration

HSRP is configured on the **interface** that is *accepting* traffic from hosts.

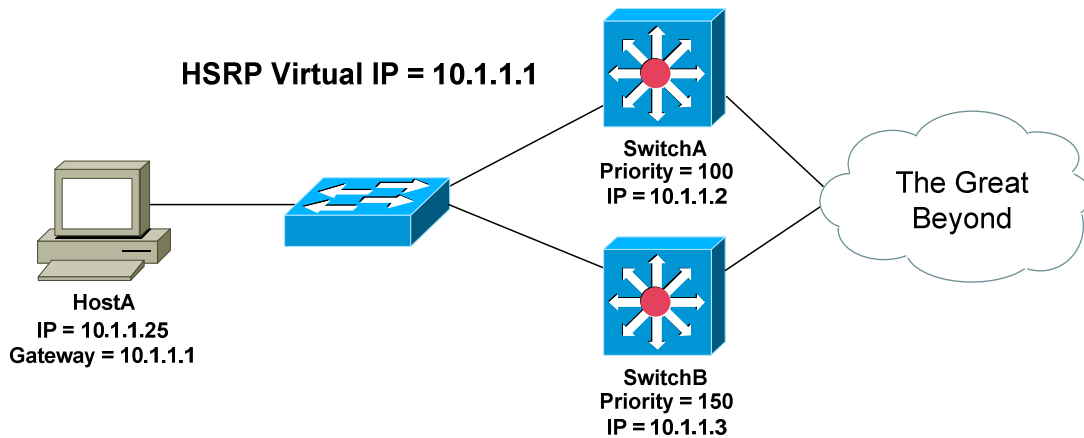
Recall that the interface with the *highest* priority is elected the active router. To configure the priority of a router from its default of **100**:

```
Router(config)# interface gi0/3
Router(config-if)# standby 1 priority 150
```

The *standby 1* command specifies the **HSRP group** the interface belongs to. HSRP can also be configured on a VLAN interface on a multilayer switch:

```
SwitchB(config)# interface vlan 100
SwitchB(config-if)# standby 1 priority 150
```

Each interface in the HSRP group retains its local IP address. The HSRP group itself is assigned a **virtual IP address**, which hosts will use as their default gateway:



To configure the virtual HSRP IP address:

```
SwitchA(config)# interface vlan 100
SwitchA(config-if)# ip address 10.1.1.2 255.255.255.0
SwitchA(config-if)# standby 1 ip 10.1.1.1
```

```
SwitchB(config)# interface vlan 100
SwitchB(config-if)# ip address 10.1.1.3 255.255.255.0
SwitchB(config-if)# standby 1 ip 10.1.1.1
```

Remember: while each multilayer switch is configured with its own *local* IP address, both are configured with the same *virtual* IP address. HostA will use this virtual IP address as its default gateway.

* * *

All original material copyright © 2014 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

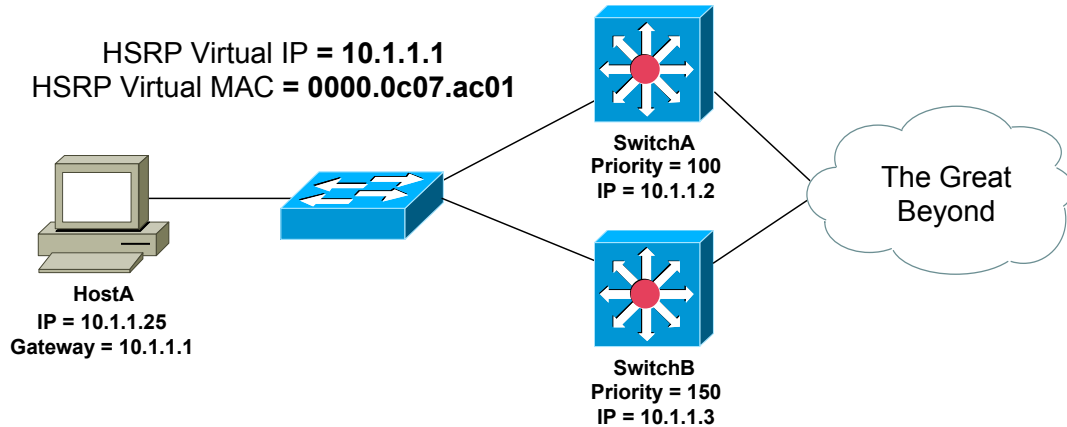
This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

HSRP Basic Configuration (continued)

HSRP supports using multiple virtual IP addresses:

```
SwitchB(config-if)# standby 1 ip 10.1.1.1
SwitchB(config-if)# standby 1 ip 10.1.1.5 secondary
```

The active router will respond to ARP requests for the virtual IP with the **virtual MAC address**.



The virtual MAC is a reserved address in the following format:

0000.0c07.acxx

...with *xx* representing the HSRP group number in hexadecimal. For example, if the HSRP group number is 8, the resulting virtual MAC address would be *0000.0c07.ac08*.

The HSRP virtual MAC address can be manually changed:

```
Switch(config-if)# standby 1 mac-address 0000.00ab.12ef
```

HSRP authentication prevents an unauthorized router from joining the HSRP group. All routers in the HSRP group must be configured with an identical authentication string.

To specify a **clear-text** authentication string:

```
Switch(config-if)# standby 1 authentication STAYOUT
```

To specify an MD5-hashed authentication string:

```
Switch(config-if)# standby 1 authentication md5 key-string 7 STAYOUT
```

* * *


All original material copyright © 2014 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

HSRP Preempt

If a new router is added to an HSRP group, it will not *preemptively* assume the role of the active router, even if it has the best priority.

In fact, the router that is first powered on will become the active router, even if it has the lowest priority!



Consider the above example:

1. If SwitchB was powered on first, it would become the active router.
2. SwitchA would be elected the standby router.
3. If SwitchB fails, SwitchA would take over as the active router.
4. Once SwitchB recovers, it *will not* retake its role of active router, *despite* having a higher priority.

The *preempt* parameter will allow a router to forcibly assume the role of active router, if it has the highest priority. The *preempt* feature is **disabled** by default:

```
SwitchB(config-if)# standby 1 preempt
```

The optional *delay* parameter will force a router to wait before preempting as the active router. The delay is measured in *seconds*:

```
Switch(config-if)# standby 1 preempt delay 10
```

The router can also be forced to wait a specified number of seconds after a *reload* before preempting the active role:

```
Switch(config-if)# standby 1 preempt reload 20
```

This allows routing protocols to converge before the router becomes active.

All original material copyright © 2014 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

HSRP Timers

Hello packets are used to elect the active and standby router, and to detect if there is a failure. By default, hello packets are exchanged every **3 seconds**.

HSRP Hello packets are sent to the multicast address **224.0.0.2** over UDP port **1985**. If no elections are occurring, only the active and standby routers exchange hello packets.

If no hello packets are received from the active router within the **holddown timer** duration, the standby router will assume it failed and take over as active. By default, the holddown timer is three times the hello timer, or **10 seconds**. Cisco's math, not mine.

To manually adjust the HSRP timers, measured in seconds:

```
SwitchB(config-if)# standby 1 timers 4 12
```

The first timer value represents the *hello* timer, while the second represents the *holddown* timer. The timers can also be specified in milliseconds:

```
SwitchB(config-if)# standby 1 timers msec 800 msec 2400
```

Troubleshooting HSRP

To view the status of each HSRP group:

```
SwitchB# show standby
```

```
VLAN 100 - Group 1
  State is Active
    1 state changes, last state change 00:02:19
  Virtual IP address is 10.1.1.1
  Active virtual MAC address is 0000.0c07.ac01
    Local virtual MAC address is 0000.0c07.ac01 (bia)
  Hello time 3 sec, hold time 10 sec
    Next hello sent in 1.412 secs
  Preemption enabled, min delay 50 sec, sync delay 40 sec
  Active router is local
  Standby router is 10.1.1.2, priority 100 (expires in 6.158 sec)
  Priority 150 (configured 150)
  Tracking 0 objects, 0 up
```

To view a more abbreviated version of this output:

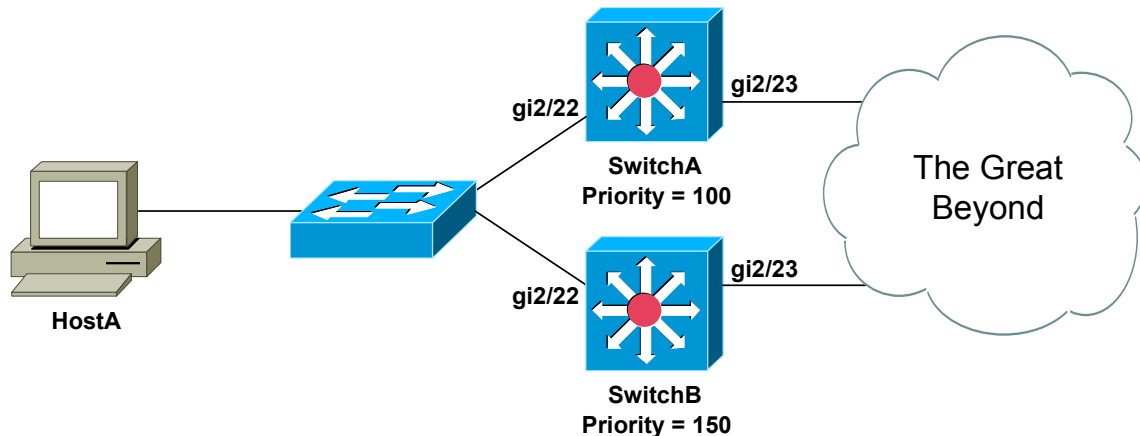
```
SwitchB# show standby brief
```

```
P indicates configured to preempt.
|
Interface Grp Prio P State Active addr Standby addr Group addr
Vlan100 1 150 P Active local 10.1.1.2 10.1.1.1
```

All original material copyright © 2014 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

HSRP Tracking



In the above example, SwitchB will become the active router, and SwitchA the standby. Both SwitchA and SwitchB exchange periodic hello packets to update their status.

If interface gi2/23 goes down on SwitchB, hello packets can still be exchanged with SwitchA via interface gi2/22. SwitchA is unaware that SwitchB has a failure and can no longer forward traffic to other networks. SwitchB will remain as the active router, and traffic will be blackholed.

To mitigate a scenario like this, HSRP can **track** interfaces. If a tracked interface fails, the router's priority is *decreased* by a specified value – by default, this is **10**.

Consider the following *tracking* configuration on SwitchB:

```
SwitchB(config-if)# standby 1 track gi2/23 60
```

If interface gi2/23 on SwitchB fails, the priority of the switch will decrease by 60. The objective is to decrement the priority enough to allow the standby router to take over as active.

This requires conscientious planning - if SwitchB's priority decremented by only 40, it would remain as active, as its priority would still be higher than SwitchA.

For tracking to be successful, the standby router must be configured to *preempt*:

```
SwitchA(config-if)# standby 1 preempt
```

Otherwise, the standby router will never take over as active.

All original material copyright © 2014 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

Virtual Router Redundancy Protocol (VRRP)

The **Virtual Router Redundancy Protocol (VRRP)** is an industry-standard Layer-3 redundancy protocol, originally defined in RFC 2338. VRRP is nearly identical to HSRP, with some notable exceptions:

- The router with the *highest* priority becomes the **master router**.
- All other routers become **backup routers**.
- The virtual MAC address is the reserved 0000.5e00.01xx, with xx representing the hexadecimal group number.
- Hello packets are sent every **1 second**, by default, and sent to multicast address 224.0.0.18.
- VRRP will *preempt* by default.
- VRRP cannot directly track interfaces – it can track an *object* which is tied to an interface, though.

Configuration of VRRP is also very similar to HSRP:

```
SwitchB(config)# interface vlan 100
SwitchB(config-if)# ip address 10.1.1.3 255.255.255.0
SwitchB(config-if)# vrrp 1 priority 150
SwitchB(config-if)# vrrp 1 authentication STAYOUT
SwitchB(config-if)# vrrp 1 ip 10.1.1.1
```

As with HSRP, the default VRRP priority is **100**, and a *higher* priority is preferred. Unlike HSRP, preemption is *enabled* by default.

To manually disable preempt:

```
Switch(config-if)# no vrrp 1 preempt
```

To view the status of each VRRP group:

```
Switch# show vrrp

Vlan 100 - Group 1
State is Master
Virtual IP address is 10.1.1.1
Virtual MAC address is 0000.5e00.0101
Advertisement interval is 3.000 sec
Preemption is enabled
min delay is 0.000 sec
Priority 150
Master Router is 10.1.1.3 (local), priority is 150
Master Advertisement interval is 3.000 sec
Master Down interval is 9.711 sec
```

All original material copyright © 2014 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

HSRP and VRRP Pseudo Load Balancing

While HSRP and VRRP do provide redundant gateways for *fault tolerance*, neither provide *load balancing* between those gateways.

Cisco will claim that load balancing is possible. Two separate HSRP or VRRP groups can be configured on a single interface:

```

SwitchA(config)# interface vlan 100
SwitchA(config-if)# ip address 10.1.1.2 255.255.255.0

SwitchA(config-if)# standby 1 priority 150
SwitchA(config-if)# standby 1 preempt
SwitchA(config-if)# standby 1 ip 10.1.1.1

SwitchA(config-if)# standby 2 priority 50
SwitchA(config-if)# standby 2 preempt
SwitchA(config-if)# standby 2 ip 10.1.1.254

SwitchB(config)# interface vlan 100
SwitchB(config-if)# ip address 10.1.1.3 255.255.255.0

SwitchB(config-if)# standby 1 priority 50
SwitchB(config-if)# standby 1 preempt
SwitchB(config-if)# standby 1 ip 10.1.1.1

SwitchB(config-if)# standby 2 priority 150
SwitchB(config-if)# standby 2 preempt
SwitchB(config-if)# standby 2 ip 10.1.1.254

```

In the above configuration, each HSRP group has been assigned a unique virtual IP address – *10.1.1.1* and *10.1.1.254* respectively. By manipulating the priority, each multilayer switch will become the *active* router for one HSRP group, and the *standby* router for the other group.

SwitchA# show standby brief

Interface	Grp	Prio	P	State	Active addr	Standby addr	Group addr
Vlan100	1	150	P	Active	local	10.1.1.3	10.1.1.1
Vlan100	2	50	P	Standby	10.1.1.3	local	10.1.1.254

SwitchB# show standby brief

Interface	Grp	Prio	P	State	Active addr	Standby addr	Group addr
Vlan100	1	50	P	Standby	10.1.1.2	local	10.1.1.1
Vlan100	2	150	P	Active	local	10.1.1.2	10.1.1.254

To achieve load balancing with this HSRP setup, half of the hosts in VLAN 100 must point to the first virtual address as their gateway – *10.1.1.1*. The other half must use the *other* virtual address as their gateway – *10.1.1.254*.

Simple and practical, right? Certainly, it is perfectly feasible to *manually* configure half of the hosts to use one gateway, and *manually* configure the other half to use another.

But hey – it's not a limitation, it's a *feature*!

<unnecessary obscene commentary edited out>

All original material copyright © 2014 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.


This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

Gateway Load Balancing Protocol (GLBP)

To overcome the shortcomings in HSRP and VRRP, Cisco developed the proprietary **Gateway Load Balancing Protocol (GLBP)**.

Routers are added to a **GLBP group**, numbered 0 to 1023. Unlike HSRP and VRRP, *multiple* GLBP routers can be active, achieving *both* redundancy and load balancing.

A priority is assigned to each GLBP interface - **100** by default. The interface with the highest priority becomes the **Active Virtual Gateway (AVG)**. If priorities are equal, the interface with the highest IP will become the AVG.



Routers in the GLBP group are assigned a single virtual IP address. Hosts will use this virtual address as their default gateway. The AVG will respond to ARP requests for the virtual IP with the virtual MAC address of an **Active Virtual Forwarder (AVF)**.

Up to three routers can be elected as AVFs. The AVG assigns a virtual MAC address to each AVF, *and* to itself, for a maximum total of **4 virtual MAC addresses**. Only the AVG and AVFs can forward traffic for hosts.

Any router not elected as an AVF or AVG will become a **Secondary Virtual Forwarder (SVF)**, and will wait in standby until an AVF fails.

* * *

All original material copyright © 2014 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

Gateway Load Balancing Protocol (GLBP) (continued)

Basic GLBP configuration is nearly identical to HSRP:

```
SwitchB(config)# interface gi2/22
SwitchB(config-if)# glbp 1 priority 150
SwitchB(config-if)# glbp 1 preempt
SwitchB(config-if)# glbp 1 ip 10.1.1.1
```

Remember that the interface with the *highest* priority is elected as the AVG. Preemption is **disabled** by default with GLBP.

What determines whether a router becomes an AVF or SVF? Each router is assigned a **weight**, and the default weight is **100**. A *higher* weight is preferred.

Weight can be configured two ways:

- Statically
- Dynamically, by tracking an interface

To manually specify the weight of a router:

```
SwitchB(config)# interface gi2/22
SwitchB(config-if)# glbp 1 weighting 150
```

Like VRRP, GLBP cannot directly *track* an interface. Instead, an interface must be specified as a **tracked object**:

```
SwitchB(config)# track 10 interface gi2/23
```

The tracked object can then be referenced as part of a *weighting* command:

```
SwitchB(config)# interface gi2/22
SwitchB(config-if)# glbp 1 weighting track 10 decrement 65
```

Note that the tracking *object-number* must match. By default, the weight will decrement by 10.

Weight *thresholds* can be configured, forcing a router to relinquish its AVF role if it falls below the minimum threshold:

```
SwitchB(config)# interface gi2/22
SwitchB(config-if)# glbp 1 weighting 150 lower 90 upper 125
```

If the weight falls below the *lower* threshold, the router must stop functioning as an AVF. The router will become an AVF again once its weight reaches the *upper* threshold, as long as *preempt* is configured.

* * *

All original material copyright © 2014 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

Gateway Load Balancing Protocol (GLBP) (continued)

The AVG assigns a virtual MAC address to itself and up to three AVFs. The AVG will respond to ARP requests for the virtual IP address with one of these virtual MAC addresses. This allows GLBP to provide load balancing in addition to redundancy.

GLBP supports three load balancing methods:

- **Round Robin**
- **Weighted**
- **Host-dependent**

The *default* load balancing method is per-host **round robin**. Traffic from hosts is distributed *equally* across all routers in the GLBP group. The AVG will respond to the first host ARP request with the first virtual MAC address. The second ARP request will receive the second virtual MAC address, etc.

The **weighted** load balancing method will distribute traffic proportionally, based on a router's *weight*. Routers with a higher weight will receive a proportionally higher percentage of traffic.

Host-dependent load balancing will provide a host device with the *same* virtual MAC address every time it performs an ARP request.

The load balancing method should be configured on the AVG:

```
SwitchB(config-if)# glbp 1 load-balancing round-robin
SwitchB(config-if)# glbp 1 load-balancing weighted
SwitchB(config-if)# glbp 1 load-balancing host-dependent
```

Other fun facts about GLBP:

- Hello packets are sent every **3 seconds**.
- Hello packets are sent to multicast address **224.0.0.102**.
- The default holddown time is **10 seconds**.
- The virtual MAC address is the reserved 0007.b4xx.xxyy, with xxxx representing the GLBP group number, and yy representing the AVF number.

To view the status of each GLBP group:

```
SwitchB# show glbp
```

* * *

All original material copyright © 2014 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

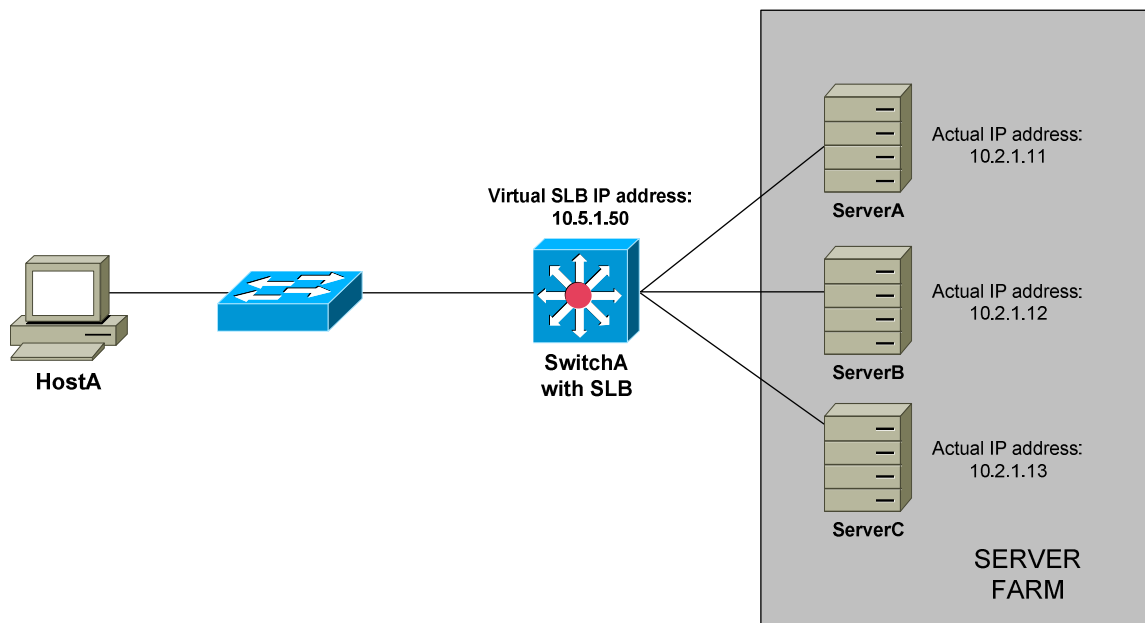
Server Load Balancing (SLB)

HSRP, VRRP, and GLBP provide *gateway* redundancy for clients. A limited number of Cisco platforms also support a basic *clustering* service.

Server Load Balancing (SLB) allows a router to apply a virtual IP address to a group of servers. Each server in the cluster should provide the same function and be configured identically, except for their IP addresses.

Hosts make requests to a single virtual IP address to access the server farm. If an individual server fails, the server farm will stay functional, as long as at least one server is online. This is also useful for seamless server repairs and maintenance.

The following diagram demonstrates SLB:



Three servers exist in the server farm. SwitchA is configured with a virtual SLB address of *10.5.1.50*. If HostA wants to access any service from the server farm, it will forward the request to virtual SLB address.

SwitchA accepts the request on behalf on the server farm, and redirects the request to one of the physical servers in the farm. This is **transparent** to the host – it is unaware of which physical server it is truly connecting to.

Note: Cisco is discontinuing support for server load-balancing mechanisms like SLB. There are many competitors that make superior products, such as F5 load balancers.

* * *

All original material copyright © 2014 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

Server Load Balancing (SLB) (continued)

SLB supports two load balancing methods:

- **Weighted round robin**
- **Weighted least connections**

Weighted round robin is the default method. Traffic is forwarded to servers in a round robin fashion. However, servers with a higher *weight* will receive a proportionally higher percentage of traffic.

In contrast, **weighted least connections** will allocate traffic based on *current usage*, and will allocate traffic to the server with the least amount of current connections.

SLB requires configuration of two separate components:

- **Server farm**
- **Virtual server**

To configure the Server Farm:

```
SwitchA(config)# ip slb serverfarm MYFARM
SwitchA(config-slb-sfarm)# predictor leastconns

SwitchA(config-slb-sfarm)# real 10.2.1.11
SwitchA(config-slb-real)# weight 150
SwitchA(config-slb-real)# inservice

SwitchA(config-slb-sfarm)# real 10.2.1.12
SwitchA(config-slb-real)# weight 100
SwitchA(config-slb-real)# inservice

SwitchA(config-slb-sfarm)# real 10.2.1.13
SwitchA(config-slb-real)# weight 75
SwitchA(config-slb-real)# inservice
```

The *ip slb serverfarm* command sets the server farm name, and enters SLB server farm configuration mode. The *predictor* command sets the load balancing method.

The *real* command identifies the IP address of a physical server in the farm. The *weight* command assigns the load balancing weight for that server. The *inservice* command activates the real server. To deactivate a specific server:

```
SwitchA(config-slb-sfarm)# real 10.2.1.12
SwitchA(config-slb-real)# no inservice
```

* * *

All original material copyright © 2014 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

SLB Configuration (continued)

To configure the virtual server:

```
SwitchA(config)# ip slb vserver VSERVERNAME
SwitchA(config-slb-vserver)# serverfarm MYFARM

SwitchA(config-slb-vserver)# virtual 10.5.1.50
SwitchA(config-slb-vserver)# client 10.1.0.0 0.0.255.255
SwitchA(config-slb-vserver)# inservice
```

The *ip slb vserver* command sets the virtual server name, and enters SLB virtual server configuration mode. The *serverfarm* command associates the server farm to this virtual server.

The *virtual* command assigns the virtual IP address for the server farm.

The *client* command specifies which hosts can access the server farm. It utilizes a wildcard mask like an access-list.

The *inservice* activates the virtual server. To deactivate the virtual server:

```
SwitchA(config-slb-vserver)# no inservice
```

To view the status of SLB server farms and virtual servers:

```
SwitchA# show ip slb serverfarms
SwitchA# show ip slb vserver
SwitchA# show ip slb real
```

* * *

All original material copyright © 2014 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.

Switch Chassis Redundancy

Modular Cisco switches support multiple **supervisor engines** for redundancy. This redundancy can be configured in one of three **modes**:

- **Route processor redundancy (RPR)**
- **Route processor redundancy plus (RPR+)**
- **Stateful switchover (SSO)**

With **route processor redundancy (RPR)**, the redundant supervisor engine is not fully initialized. If the primary supervisor fails, the standby supervisor must reinitialize all other switch modules in the chassis before functionality is restored. This results in an outage that lasts several minutes.

With **route processor redundancy plus (RPR+)**, the redundant supervisor engine is fully initialized, but does not participate in or cache any Layer-2 or Layer-3 functions.

If the primary supervisor fails, the standby supervisor can immediately activate all Layer-2 and Layer-3 functions, without having to reinitialize all other switch modules in the chassis. This process usually takes less than a minute.

With **stateful switchover (SSO)**, the redundant supervisor engine is fully initialized, and all Layer-2 and Layer-3 functions are synchronized with the primary supervisor. If the primary supervisor fails, failover to the standby supervisor will occur *immediately*.

To enable redundancy on the switch, and to choose the appropriate redundancy mode:

```
Switch(config)# redundancy
Switch(config-red)# mode rpr
Switch(config-red)# mode rpr-plus
Switch(config-red)# mode sso
```

The redundancy commands must be configured identically on *both* supervisor engines.

Note: RPR+ mode requires that both supervisor engines utilize the *exact* same version of the Cisco IOS.

* * *

All original material copyright © 2014 by Aaron Balchunas (aaron@routeralley.com), unless otherwise noted. All other material copyright © of their respective owners.

This material may be copied and used freely, but may not be altered or sold without the expressed written consent of the owner of the above copyright. Updated material may be found at <http://www.routeralley.com>.