- Redundancy and Load Balancing -

Importance of Redundancy

Consider the following example:

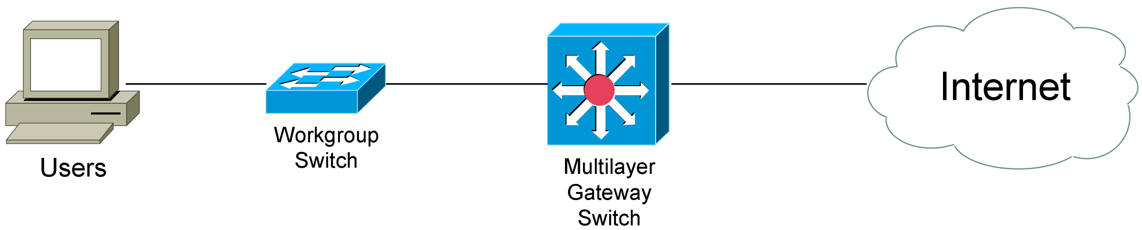

The users utilize a single *gateway* to reach the Internet. In this example, the gateway is a multilayer switch; however, a Layer-3 router is just as common. Throughout the rest of this section, the terms *router* and *multilayer switch* will be used interchangeably.

The gateway represents a single point of failure on this network. If that gateway fails, users will lose access to all resources beyond that gateway. This lack of redundancy may be unacceptable on business-critical systems that require maximum uptime.

It is possible to provide multiple gateways for host devices:

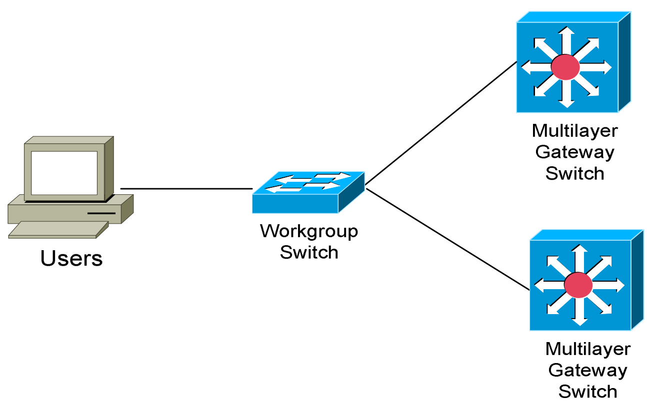

However, this required a solution transparent to the end user (or host device). Cisco devices support three protocols that provide this transparent redundancy:

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

Hot Standby Router Protocol (HSRP)

Cisco developed a proprietary protocol named **Hot Standby Router Protocol (HSRP)** that allows multiple routers or multilayer switches to masquerade as a single gateway. This is accomplished by assigning a **virtual IP address** to all routers participating in HSRP.

All routers are assigned to a single HSRP **group** (numbered 0-255). Note however, that most Catalyst switches will support only **16** configured HSRP groups. HSRP routers are elected to specific roles:

- **Active Router –** the router currently serving as the gateway.
- **Standby Router –** the 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. HSRP routers regularly send Hello packets (by default, every **3 seconds)** to ensure all routers are functioning. If the current Active Router fails, the Standby Router is made active, and a new Standby is elected.

The role of an HSRP router is dictated by its **priority**. The priority can range from 0 – 255, with a default of **100.** The router with the **highest** (a higher value is *better*) **priority** is elected the Active Router; the router with the **second highest priority** becomes the Standby Router. If all priorities are equal, whichever router has the **highest IP Address** on its HSRP interface is elected the Active Router.

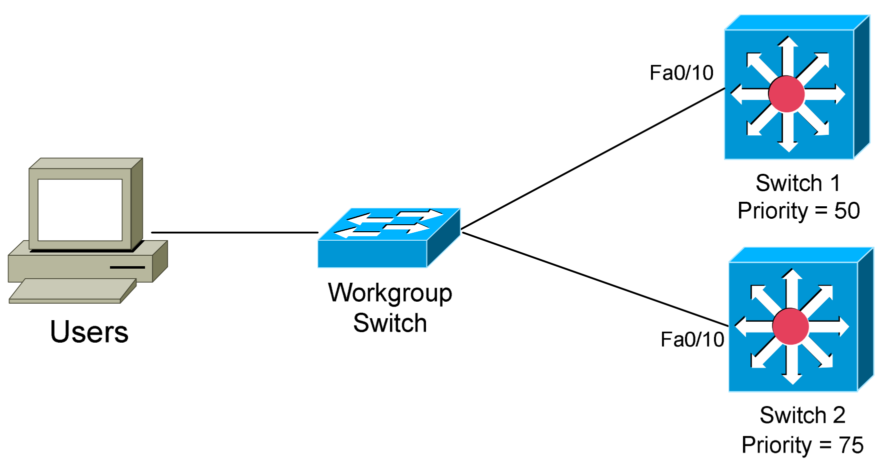

In the above example, Switch 2 would become the Active HSRP router, as it has the highest priority. Switch 1 would become the Standby router.

HSRP States

A router or multilayer switch configured for HSRP will progress through several **states** before settling into a role:

- **Disabled –** the interfaces is not configured for HSRP, or is administratively shut down.
- **Init –** this is the *starting* state when an interface is first brought up.
- **Learn** – the router is waiting to hear hellos from the Active Router, to learn the configured Virtual Address.
- **Listen –** the router has learned the Virtual IP address, but was not elected the Active or Standby Router.
- **Speak –** the router is currently participating in an Active Router election, and is sending Hello packets.
- **Standby –** the router is acting as a backup to the Active Router. Standby routers monitor and send hellos to the Active Router.
- **Active –** the router is currently accepting and forwarding user traffic, using the Virtual IP address. The Active Router actively exchanges hellos with the Standby Router.

By default, HSRP Hello packets are sent **every 3 seconds**.

Routers in a **listening** state will only *listen for* and *not periodically send* hello packets. While the HSRP is fully converged, only the **Active** and **Standby** Routers will send hellos. Routers will also send out hellos when **Speaking**, or electing the Active and Standby routers.

When electing the Active and Standby routers, the routers will enter a **Speaking** state. HSRP hellos are used to complete the election process.

Thus, the three states which send out hello packets as follows:
- **Speak**
- **Standby**
- **Active**

HSRP Configuration

All HSRP configuration is completed on the **interface** that is *accepting* traffic on behalf of host devices.

To configure the priority of a router:

Switch(config)# *interface fa0/10*
Switch(config-if)# *standby 1 priority 150*

The *standby 1* command specifies the HSRP group that interface belongs to. The *priority 150* parameter changes the actual priority value. Remember that a higher value is preferred, and that the default priority is **100**.

However, if a new router is added to the HSRP group, and it has the best priority, it will *not* automatically assume the role of the Active router. In fact, the first router to be powered on will become the Active router, even if it has the lowest priority!

To force the highest-priority router to assume the role of Active router:

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

The *standby 1 preempt* command allows this switch to force itself as the Active router, if it has the highest priority. The optional *delay 10* parameter instructs the router to wait 10 seconds before assuming an Active status.

HSRP routers send out Hello packets to verify each other's status:

Switch(config-if)# *standby 1 timers 4 12*

The *standby 1 timers* command configures the two HSRP timers. The first setting *4* sets the Hello timer to 4 seconds. The second setting *12* sets the holddown timer to 12 seconds.

Remember, by default, Hello packets are sent **every 3 seconds**. Only the Standby router listens to Hello packets from the Active router. If the Standby router does not hear any Hellos from the Active router for the holddown period, then it will assume the Active router is down.

In general, the holddown timer should be three times the Hello timer (the default holddown time is **10 seconds**). HSRP Hello packets are sent to the multicast address **224.0.0.2** over UDP port **1985**.

(Reference: http://www.cisco.com/en/US/docs/internetworking/case/studies/cs009.html)

HSRP Configuration (continued)

Each router in the HSRP group retains the address configured on its local interface. However, the HSRP group itself is assigned a **virtual IP address**. Host devices use this virtual address as their default gateway.

To configure the virtual HSRP IP address:

Switch(config)# *int fa0/10*
Switch(config-if)# *standby 1 ip 192.168.1.5*

Multiple virtual HSRP IP addresses can be used:

Switch(config-if)# *standby 1 ip 192.168.1.5*
Switch(config-if)# *standby 1 ip 192.168.1.6 secondary*

The HSRP group is also assigned a virtual MAC address. By default, a reserved MAC address is used:

0000.0c07.ac*xx*

…where *xx* is the HSRP group number in hexadecimal. For example, if the HSRP Group number was 8, the resulting virtual MAC address would be:

0000.0c07.ac08

The HSRP virtual MAC address can be manually specified:

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

Authentication can be configured for HSRP. All HSRP routers in the group must be configured with the same authentication string. To specify a **clear-text** authentication string**:**

Switch(config-if)# *standby 1 authentication CISCO*

To specify an MD5-hashed authentication string:

Switch(config-if)# *standby 1 authentication md5 key-string 7 CISCO*

HSRP Tracking

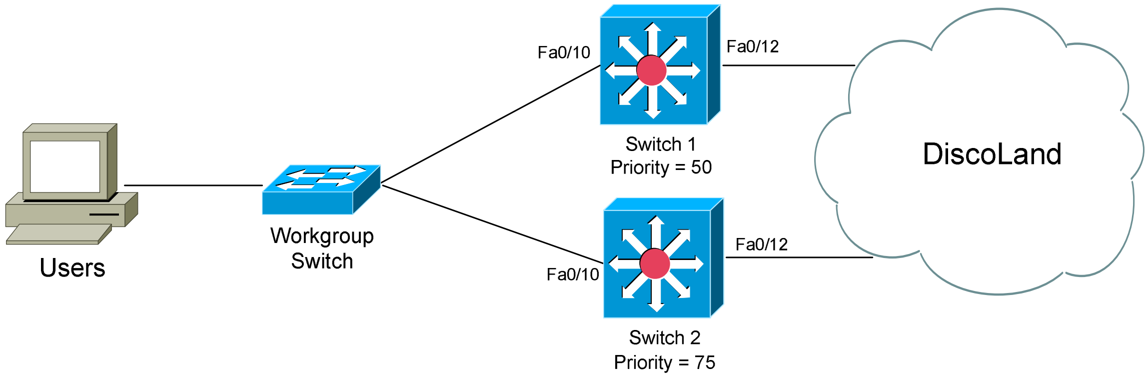

In the above example, Switch 2 becomes the Active Router, and Switch 1 becomes the Standby router. Both Switch 1 and Switch 2 send out Hello packets with updates on their status.

On Switch 2, if port Fa0/12 goes down, the switch is still able to send Hello packets to Switch 1 via Fa0/10. Thus, Switch 1 is unaware that Switch 2 is no longer capable of forwarding traffic, as Switch 2 still appears to be *active* (sending hellos).

To combat this, HSRP can **track** interfaces. If the tracked interface fails, the router's (or multilayer switch's) priority is *decreased* by a specific value. Observe the following tracking configuration on Switch 2:

Switch2(config-if)# *standby 1 track fa0/12 50*

The above command sets tracking for the *fa0/12* interface, and will decrease the priority of the switch by *50* if the interface fails. The objective is to decrement the priority enough to allow another router to assume an Active status. This requires conscientious planning by the network administrator. In the above example, Switch 2's priority would be decremented to 25 if its fa0/12 interface failed, which is less than Switch 1's priority of 50.

Tracking of interfaces will not be successful unless the other router is configured to **preempt** the current Active Router.

Switch1(config-if)# *standby 1 preempt*

If the above command was not present, Switch 1 would never assume an Active state, even if Switch 2's priority was decreased to 1.

Practical HSRP Example

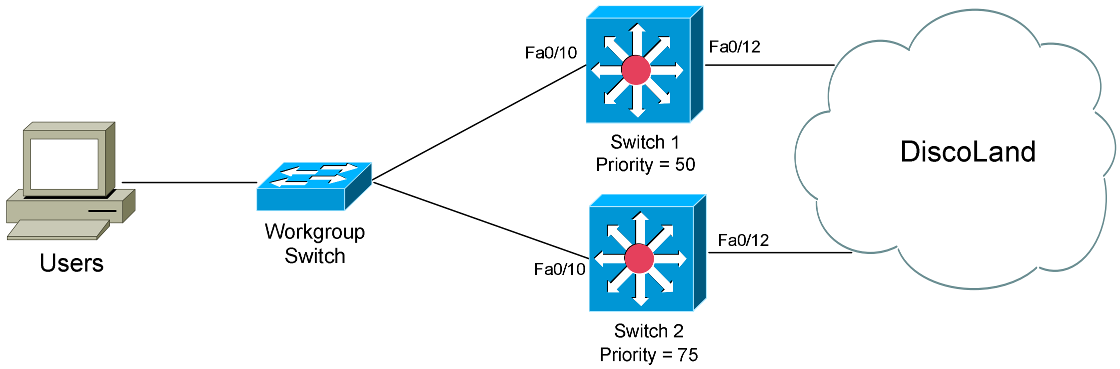

Switch1(config)# *int fa0/10*
Switch1(config-if)# *no switchport*
Switch1(config-if)# *ip address 192.168.1.5 255.255.255.0*
Switch1(config-if)# *standby 1 priority 50*
Switch1(config-if)# *standby 1 preempt*
Switch1(config-if)# *standby 1 ip 192.168.1.1*
Switch1(config-if)# *standby 1 authentication CISCO*

Switch2(config)# *int fa0/10*
Switch2(config-if)# *no switchport*
Switch2(config-if)# *ip address 192.168.1.6 255.255.255.0*
Switch2(config-if)# *standby 1 priority 75*
Switch2(config-if)# *standby 1 preempt*
Switch2(config-if)# *standby 1 ip 192.168.1.1*
Switch2(config-if)# *standby 1 authentication CISCO*
Switch2(config-if)# *standby 1 track fa0/12 50*

The *no switchport* command specifies that interface *fa0/10* is a Layer-3 (routed) port. Both switches are assigned a unique *ip address* to their local interfaces. However, both are given a single **HSRP virtual IP** address. Host devices will use this virtual address as their default gateway.

Because of its higher priority, Switch 2 will become the Active Router. Its priority will decrement by 50 if interface fa0/12 should fail. Because Switch 1 is configured with the *preempt* command, it will take over as the Active Router if this should occur.

To view the status of a configured HSRP group:

Switch2# *show standby*

```
Fastethernet0/10 - Group 1
  State is Active
    1 state changes, last state change 00:02:19
  Virtual IP address is 192.168.1.1
  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 192.168.1.5, priority 50 (expires in 6.158 sec)
  Priority 75 (configured 75)
    Tracking 1 objects, 1 up
```

Virtual Router Redundancy Protocol (VRRP)

The industry-standard equivalent of HSRP is the **Virtual Router Redundancy Protocol (VRRP)**, defined in RFC 2338. It 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.**
- By default, the virtual MAC address is 0000.5e00.01*xx*, where *xx* is the hexadecimal group number.
- Hellos are sent every **1 second**, by default.
- VRRP Hellos are sent to multicast address 224.0.0.18.
- VRRP will *preempt* by default.
- **VRRP cannot track interfaces.**

Configuration of VRRP is also very similar to HSRP:

Switch(config)# *int fa0/10*
Switch(config-if)# *no switchport*
Switch(config-if)# *ip address 192.168.1.6 255.255.255.0*
Switch(config-if)# *vrrp 1 priority 75*
Switch(config-if)# *vrrp 1 authentication CISCO*
Switch(config-if)# *vrrp 1 ip 192.168.1.1*

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 VRRP status:

Switch# *show vrrp*

```
Fastethernet 0/10 - Group 1
State is Master
Virtual IP address is 192.168.1.1
Virtual MAC address is 0000.5e00.0101
Advertisement interval is 3.000 sec
Preemption is enabled
min delay is 0.000 sec
Priority 75
Master Router is 192.168.1.6 (local), priority is 75
Master Advertisement interval is 3.000 sec
Master Down interval is 9.711 sec
```

(Reference: http://www.cisco.com/en/US/docs/ios/12_0st/12_0st18/feature/guide/st_vrrpx.html)

HSRP's and VRRP's "Pseudo" Load-Balancing

While HSRP and VRRP do provide *redundant* gateways for *fault tolerance*, they *do not* provide *load-balancing* between those gateways.

Cisco pretends that load balancing is possible. Theoretically, two separate HSRP or VRRP groups can be configured on *each* router:

Switch1(config)# *int fa0/10*
Switch1(config-if)# *no switchport*
Switch1(config-if)# *ip address 192.168.1.5 255.255.255.0*

Switch1(config-if)# *standby 1 priority 100*
Switch1(config-if)# *standby 1 preempt*
Switch1(config-if)# *standby 1 ip 192.168.1.1*

Switch1(config-if)# *standby 2 priority 50*
Switch1(config-if)# *standby 2 preempt*
Switch1(config-if)# *standby 2 ip 192.168.1.2*

Switch2(config)# *int fa0/10*
Switch2(config-if)# *no switchport*
Switch2(config-if)# *ip address 192.168.1.6 255.255.255.0*

Switch2(config-if)# *standby 1 priority 50*
Switch2(config-if)# *standby 1 preempt*
Switch2(config-if)# *standby 1 ip 192.168.1.1*

Switch2(config-if)# *standby 2 priority 100*
Switch2(config-if)# *standby 2 preempt*
Switch2(config-if)# *standby 2 ip 192.168.1.2*

In the above example, each HSRP group (*1* and *2*) has been assigned a unique virtual IP address. By adjusting the priority, each multilayer switch will become the Active router for one HSRP group, and the Standby router for the other group.

Switch1# *show standby brief*

```
Interface  Grp Prio P State    Active addr   Standby addr   Group addr
Fa0/10     1   100  P Active   local         192.168.1.6    192.168.1.1
Fa0/10     2   50   P Standby  192.168.1.6   local          192.168.1.2
```

Switch2# *show standby brief*

```
Interface  Grp Prio P State    Active addr   Standby addr   Group addr
Fa0/10     1   50   P Standby  192.168.1.5   local          192.168.1.1
Fa0/10     2   100  P Active   local         192.168.1.5    192.168.1.2
```

To achieve HSRP redundancy with this setup, half of the host devices would need to point to first virtual address (192.168.1.1), and the remaining half to the *other* virtual address (192.168.1.2).

That's simple and dynamic, right? Nothing like having to *manually* configure half of the clients to use one gateway address, and half of them to use the other. Or set up two separate DHCP scopes….

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

<unnecessary obscene commentary edited out>

Gateway Load Balancing Protocol (GLBP)

To overcome the…. *shortcomings* in HSRP and VRRP, Cisco developed the oh-so proprietary **Gateway Load Balancing Protocol (GLBP).** Routers or multilayer switches are added to a GLBP group - but unlike HSRP/VRRP, **all routers** are Active. Thus, both redundancy and load-balancing are achieved. GLBP utilizes multicast address **224.0.0.102.**

As with HSRP and VRRP, GLBP routers are placed in a group (1-255). Routers are assigned a priority (default is 100) - the router with the highest priority becomes the **Active Virtual Gateway (AVG)**. If priorities are equal, the router with the highest IP on its interface will become the AVG.

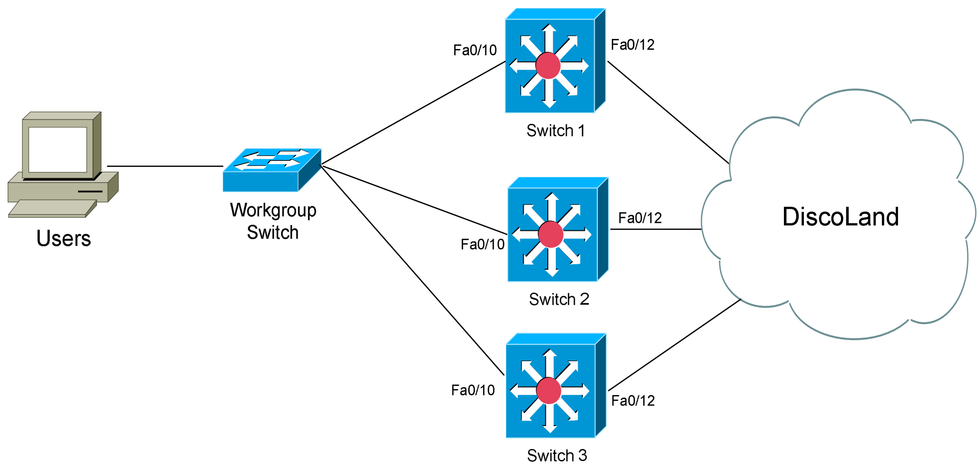

Routers in the GLBP group are assigned a single virtual IP address. Host devices will use this virtual address as their default gateway, and will broadcast an ARP request to determine the MAC address for that virtual IP. The router elected as the AVG listens for these ARP requests.

In addition to the AVG, up to three other routers can elected as **Active Virtual Forwarders (AVF's)**. The AVG assigns each AVF (including *itself*) a virtual MAC address, for a maximum total of **4 virtual MAC addresses**. When a client performs an ARP request, the AVG will provide the client with one of the virtual MAC addresses. In this way, load balancing can be achieved.

GLBP is *not* limited to four routers. Any router not elected to be an AVF will become a **Secondary Virtual Forwarder (SVF),** and will wait in standby until an AVF fails.

(Reference: http://www.cisco.com/en/US/docs/ios/12_2t/12_2t15/feature/guide/ft_glbp.html)

Gateway Load Balancing Protocol (GLBP) (continued)

What determines whether a router becomes an AVF or SVF? Each router is assigned a **weight**, and the default weight is 100. Weight can be statically configured, or dynamically decided by the router. When dynamically decided, a router's weight will drop if a **tracked** interface fails. Weight *thresholds* can be configured, forcing a router to relinquish its AVF status if it falls below the minimum threshold.

GLBP supports three load balancing methods:

- **Round Robin** – Traffic is distributed equally across all routers. The first host request receives Router 1's virtual MAC address, the second request will receive Router 2's virtual MAC address, etc. This is the *default load balancing* mechanism.
- **Weighted –** Traffic is distributed to routers proportional to their configured *weight.* Routers with a higher weight will be utilized more frequently.
- **Host-Dependent –** A host device will always receive the *same* virtual MAC-address when it performs an ARP request.

To configure a GLBP router's priority to 150, and enable preempt (preemption is not enabled by default):

Switch(config)# *int fa0/10*
Switch(config-if)# *glbp 1 priority 150*
Switch(config-if)# *glbp 1 preempt*

To track an interface, to reduce a router's weight if that interface fails:

Switch(config)# *track 10 interface fa0/12*

Switch(config-if)# *glbp 1 weighting track 10 decrement 50*

The first command creates a *track* **object** *10*, which is tracking *interface fa0/12*. The second command assigns that track object to *glbp* group *1*, and will decrease this router's weight by *50* if interface fa0/12 fails. Another router cannot become an AVF unless it is configured to *preempt*.

To specify the Virtual IP, and the load-balancing method:

Switch(config-if)# *glbp 1 ip 192.168.1.2*
Switch(config-if)# *glbp 1 load-balancing weighted*

Server Load Balancing (SLB)

HSRP, VRRP, and GLBP provide *gateway* redundancy for clients. Cisco routers and switches also support a basic *clustering* service.

Server Load Balancing (SLB) allows a router to apply a virtual IP address to a group of servers. All of the servers should be configured identically (with the exception of their IP addresses), and provide the same function. Having multiple servers allows for both redundancy and load-balancing.

Clients point to a single virtual IP address to access the server farm. The client is unaware of which server it is truly connecting to. If a specific server fails, the server farm can stay operational. Individual servers can be brought down for repair or maintenance, and the server farm can stay functional.

The following diagram demonstrates SLB:

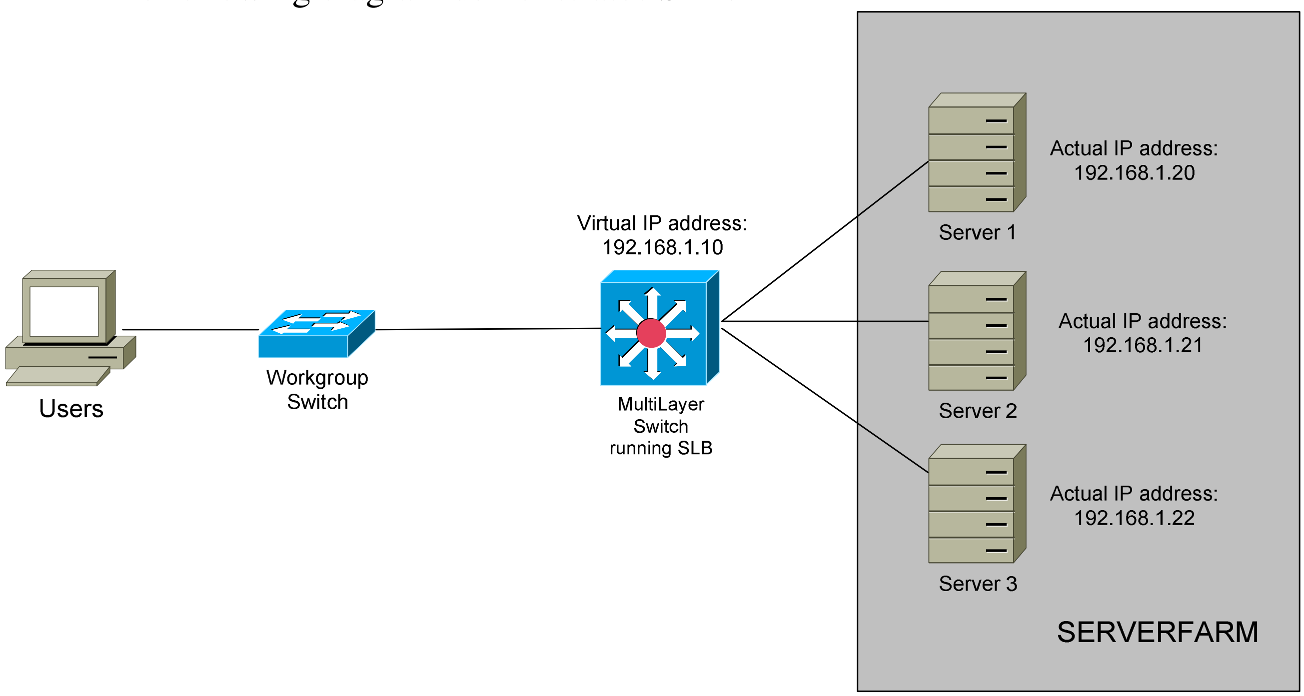

Assume the servers are Web servers. To access the Web resource, users will connect to the Virtual IP address of 192.168.1.10. The multilayer switch intercepts this packet, and redirects it to one of the physical servers inside the server farm. In essence, the multilayer switch is functioning as a Virtual Server.

SLB Load Balancing

Two load balancing methods exist for SLB:

- **Weighted Round Robin –** Traffic is forwarded to the physical servers in a round robin fashion. However, servers with a higher *weight* are assigned more traffic. This is the default method.
- **Weighted Least Connections** – Traffic is assigned to the server with the least amount of current connections.

SLB Configuration

Two separate elements need to be configured with SLB, the **Server Farm**, and the **Virtual Server**. To configure the Server Farm:

Switch(config)# *ip slb serverfarm MYFARM*
Switch(config-slb-sfarm)# *predictor leastconns*

Switch(config-slb-sfarm)# *real 192.168.1.20*
Switch(config-slb-real)# *weight 150*
Switch(config-slb-real)# *inservice*

Switch(config-slb-sfarm)# *real 192.168.1.21*
Switch(config-slb-real)# *weight 100*
Switch(config-slb-real)# *inservice*

Switch(config-slb-sfarm)# *real 192.168.1.22*
Switch(config-slb-real)# *weight 75*
Switch(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, and enters SLB Real Server configuration mode. The *weight* command assigns the load-balancing weight for that server. The *inservice* command activates the real server. To deactivate a specific server:

Switch(config-slb-sfarm)# *real 192.168.1.22*
Switch(config-slb-real)# *no inservice*

(Reference: http://www.cisco.com/en/US/docs/ios/12_1/12_1e8/feature/guide/iosslb8e.html)

SLB Configuration (continued)

To configure the Virtual Server:

Switch(config)# *ip slb vserver VSERVERNAME*
Switch(config-slb-vserver)# *serverfarm MYFARM*

Switch(config-slb-vserver)# *virtual 192.168.1.10*
Switch(config-slb-vserver)# *client 192.168.0.0 0.0.255.255*
Switch(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 clients can access the server farm. It utilizes a wildcard mask like an access-list. In the above example, *client 192.168.0.0 0.0.255.255* would allow all clients in the 192.168.x.x Class B network.

The *inservice* activates the Virtual Server. To deactivate a Virtual Server:

Switch(config-slb-vserver)# *no inservice*

To troubleshoot SLB:

Switch# *show ip slb serverfarms*
Switch# *show ip slb vserver*
Switch# *show ip slb real*

Switch Chassis Redundancy

Modular Catalyst switches support the installation of multiple Supervisor Engines for redundancy. This redundancy can be configured in one of three **modes:**

- **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 process can take several minutes.

- **Route Processor Redundancy Plus (RPR+) –** The redundant Supervisor engine is fully initialized, but performs no Layer-2 or Layer-3 functions. If the primary Supervisor fails, the standby Supervisor will activate 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.

- **Stateful Switchover (SSO) –** The redundant Supervisor engine is fully initialized, and synchronizes all Layer-2 and Layer-3 functions with the primary Supervisor. If the primary Supervisor fails, failover can occur *immediately* to the standby Supervisor.

To enable redundancy on the Catalyst 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 would need to be enabled on both Supervisor engines. RPR+ mode requires that both Supervisor engines utilize the *exact* same version of the Cisco IOS.

(Reference: http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/prod_white_paper0900aecd801c5cd7.html.
http://www.cisco.com/en/US/prod/collateral/switches/ps5718/ps708/prod_white_paper09186a0080088874.html)