



WHITE PAPER

GATEWAY LOAD BALANCING PROTOCOL OVERVIEW

This document describes the Gateway Load Balancing Protocol (GLBP) feature of Cisco IOS® Software. Cisco continues to develop the core technologies demanded by Enterprises, in response to customers' wants and needs. GLBP was developed to provide automatic, first-hop gateway load balancing, which allows for more efficient resource usage and reduced administrative costs. It is an extension of Hot Standby Router Protocol (HSRP) and specifies a protocol that dynamically assigns responsibility for a virtual IP address and distributes multiple virtual MAC addresses to members of a GLBP group.

CISCO IOS® HIGH AVAILABILITY

Leading businesses around the globe rely on the network to provide connectivity to the most mission critical applications, and many companies cannot function and even record large losses when network connections are down. By combining element redundancy, network resilience, and industry-leading experience with operational processes, Cisco can help Enterprises and Service Providers achieve 99.999% availability. GLBP is one of many Cisco IOS High Availability features. It offers enhanced redundancy and increases performance by fully utilizing the available network resources.

Like HSRP, GLBP supplies a method of providing nonstop path redundancy for IP by sharing protocol and Media Access Control (MAC) addresses between redundant gateways. Additionally, GLBP allows a group of routers or layer 3 switches to share the load of the default gateway on a Local Area Network (LAN). It therefore improves performance by facilitating better use of network resources when multiple upstream paths are available, and increases reliability and network availability by removing the single point of failure (the first hop router or layer 3 switch). GLBP enables a router to automatically assume the forwarding function of another router in the group if there is a failure in any other gateway router.

GLBP Advantages

- **Efficient use of network resources:** multiple paths upstream from the gateways can be utilized simultaneously.
- **Higher availability:** GLBP offers enhanced redundancy eliminating single point of failure of the first-hop gateway. An enhanced object-tracking feature can be used with GLBP to ensure the redundancy implementation mirrors network capabilities. This same feature is also available for HSRP and VRRP.
- **Automatic load balancing:** Off-net traffic is shared among available gateways on a per-host basis, according to the defined load-balancing algorithm.
- **Lower administration costs:** Since all hosts on a subnet can use a common default gateway while load balancing is still achieved, administration of multiple groups and gateways is unnecessary.
- **Simpler Access-layer design:** More efficient use of resources is now possible without configuring additional VLANs and subnets.

GLBP can be used if IP hosts on the LAN have a default gateway configured or learned via DHCP. It allows them to send packets to hosts on other network segments while balancing their traffic among multiple gateways.

PROBLEM DEFINITION

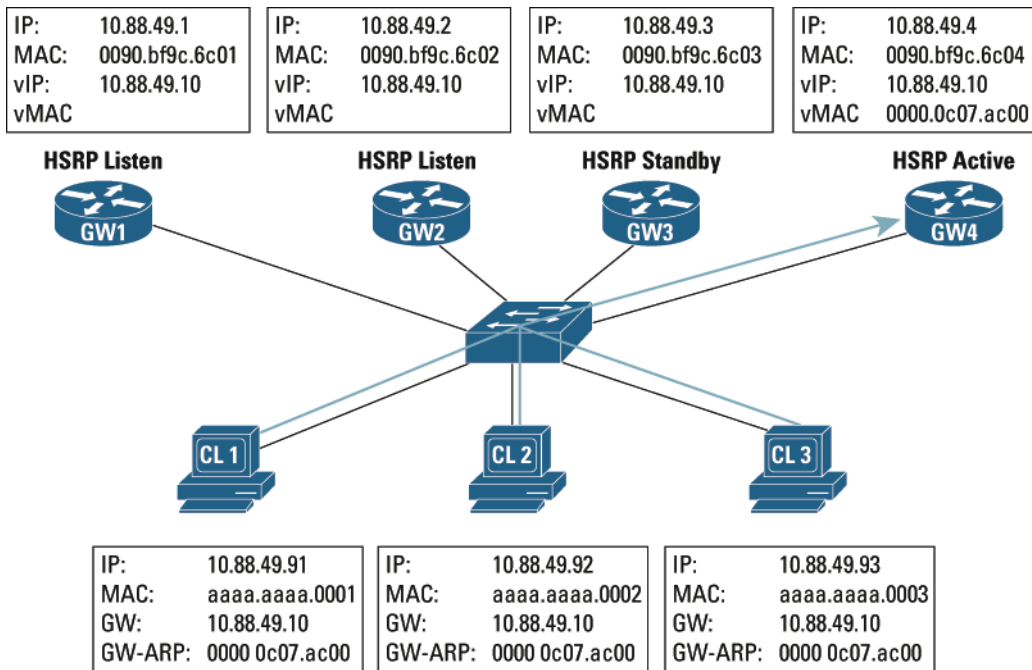
In many networks today, HSRP provides redundant default gateway services in Cisco IOS Software. While HSRP provides necessary redundancy, it does not make efficient use of all available network paths in non-fault conditions without undertaking additional configuration and design steps.

HSRP was developed and patented (Patent: 5,473,599 issued December 5, 1995) by Cisco Systems and first released in Cisco IOS Software Release 10.0. It was made available as an IETF Informational RFC (RFC 2281).

With HSRP in operation on a common IP subnet, two or more routers or layer-3 switches share a virtual IP address and a virtual MAC address to form a HSRP group. HSRP specifies an election protocol that dynamically assigns responsibility for a virtual router to one of the HSRP routers in the HSRP group. The HSRP router controlling the virtual IP address is called the *active router*. Clients on the subnet set their default gateway to the virtual IP address of the HSRP group. During normal operation, the active router responds to all ARP requests for the virtual IP address and forwards all packets sent to this IP address (Figure 1). The other members of the HSRP group remain in Standby or Listen mode and do not forward any traffic sent to the virtual IP address. **Therefore, any uplink(s) on the member routers of the HSRP group that are not forwarding traffic remains idle since there is no automatic support for load balancing in HSRP.**

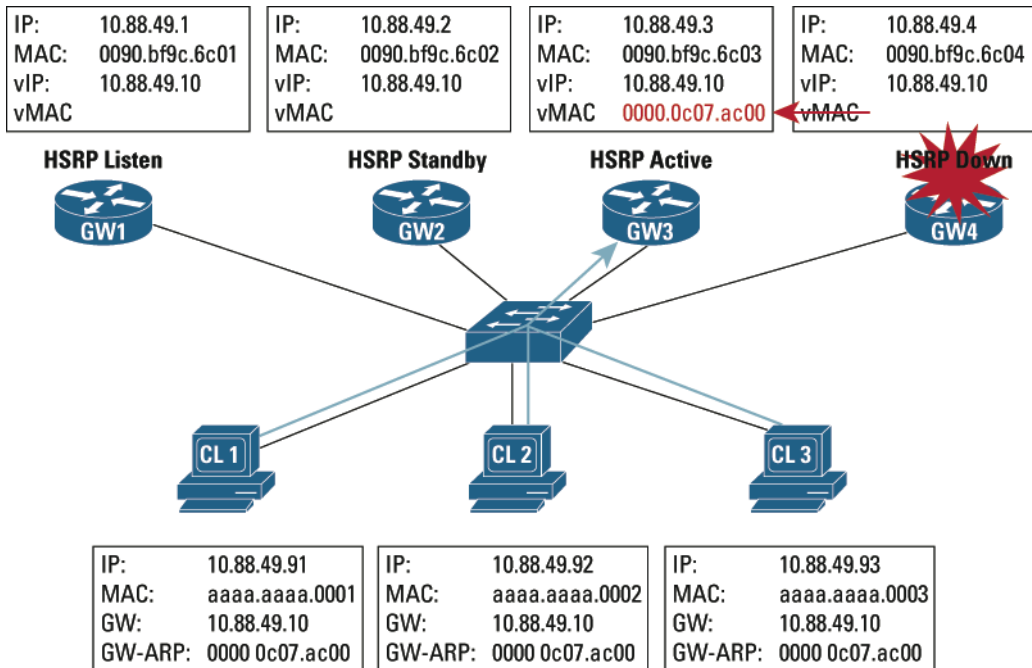
GLBP provides enhanced load balancing over HSRP.

Figure 1. HSRP in Normal Operation



HSRP provides dynamic switchover of the forwarding responsibility should the active router become unavailable. The other routers in the HSRP group detect a failure to the active router (see Figure 2). The standby router assumes the virtual MAC address and begins forwarding traffic. This process can be tuned to occur quickly so that applications and connections at the hosts are unaware that an outage has occurred.

Figure 2. HSRP in Fail-Over



HSRP does play an important role in increasing availability by negating the impact of an outage at the first-hop access point. However, it does not allow for efficient and total use of network resources when multiple available paths exist **unless additional configuration steps are taken**. The same can be said for Virtual Router Redundancy Protocol (VRRP), an IETF standard (RFC 2338) first-hop redundancy protocol that is similar to HSRP.

ALTERNATE METHODS FOR LOAD BALANCING

Many Service Providers and Enterprises require load balancing across multiple gateways and multiple uplinks for a common subnet. In many cases, these customers want to offer increased performance and take advantage of the costly, redundant facilities that have been put in place to ensure high availability during times when both the primary and redundant paths are active.


There are essentially two methods available for load balancing traffic across multiple upstream paths at the extreme edge of the network: use Layer 2 load balancing or Layer 3 load balancing with multi-group HSRP (MHSRP).

LAYER 2 LOAD BALANCING ALTERNATIVE

When load-balancing across redundant uplinks from an access layer switch in an enterprise campus or a service provider data center is required, one option is to use Layer 2 mechanisms, which include an additional spanning tree feature from Cisco called Per-VLAN-Spanning-Tree (PVST) Plus. With PVST+, two or more VLANs are created on the access layer switch. Load balancing is achieved by ‘tuning’ half of the VLANs to use the forwarding path through one uplink, and the second half to use the other uplink.

Tuning forces one distribution layer switch to be the root bridge for VLAN A, while the other distribution layer switch is the root bridge for VLAN B. Both distribution layer switches are linked with an additional trunk, which carries both VLANs. This forms a triangular loop. The spanning tree protocol (STP) removes the loop, leaving the desired forwarding arrangement. This results in balancing the load over the uplinks according to VLAN while retaining the desired redundancy.

Figure 3. Multi-VLAN Load Balancing



Fast Layer 2 spanning tree recovery is achieved using another feature from Cisco called UplinkFast. Enabling the trunks from each access switch for UplinkFast results in a recovery time of about three seconds when there is a failure in an individual uplink.

Again, this load-balancing implementation requires STP to be present in a triangular switch arrangement, whereby each access switch is homed to each distribution switch and the distribution switches are connected together by a trunk supporting all VLANs.

There are several drawbacks to this design:

- Multiple VLANs and IP subnets are required, and the user community must be divided among the multiple VLANs.
- Added complexity: campus networks were designed this way to take advantage of fast Layer 2 switching. However, the introduction of high-performance layer-3 switching in hardware negates much of the performance advantage of Layer 2 switching.

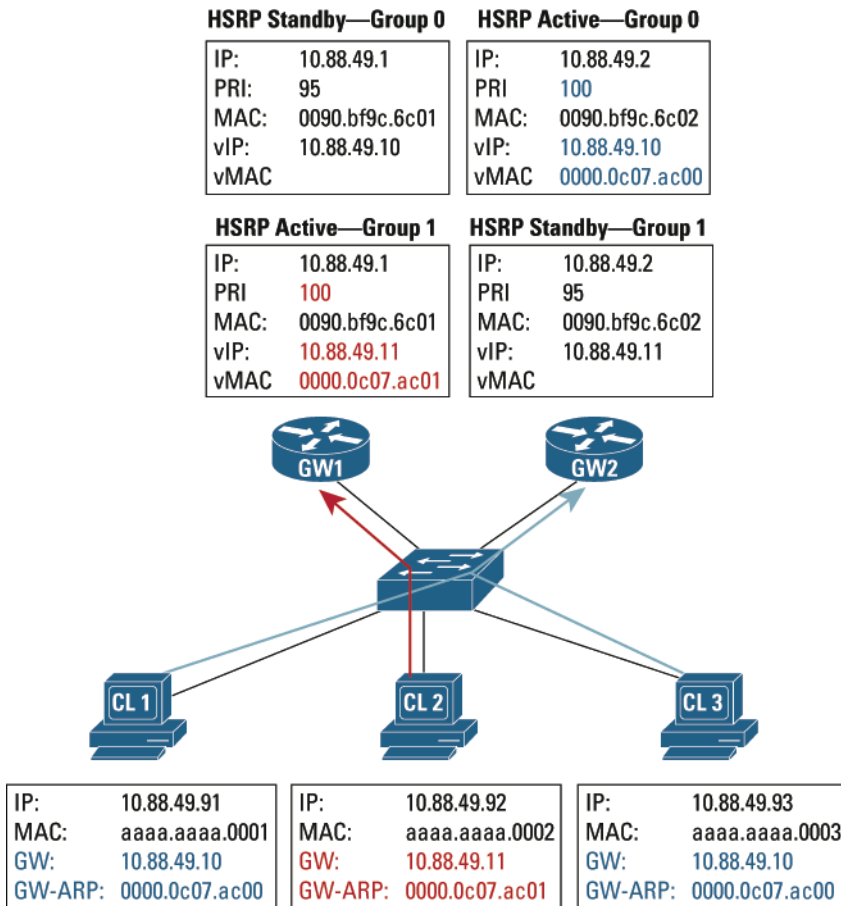
VLANs are currently best used to implement policy. Cisco Architecture for Voice, Video, and Integrated Data (AVVID) and best practice design suggests that one IP subnet should map to a single VLAN and a single switch in a wiring closet. (See *Cisco AVVID Network Infrastructure, Gigabit Campus Network Design: Principles and Architecture* at

http://www.cisco.com/en/US/netsol/ns340/ns394/ns74/ns149/networking_solutions_white_paper09186a00800a3e16.shtml.)

LAYER 3 MULTI-GROUP HSRP LOAD BALANCING

Multiple HSRP groups can be configured on a common subnet to provide load balancing across multiple gateways. A router that is “active” for one group is “standby” for the other. HSRP priority determines the initial state. The router with the highest priority will become the active router for that group. Clients on the subnet are configured with different default gateways with different virtual IP addresses corresponding to the different HSRP groups. An example is shown in Figure 4.

Figure 4. Multi-Group HSRP

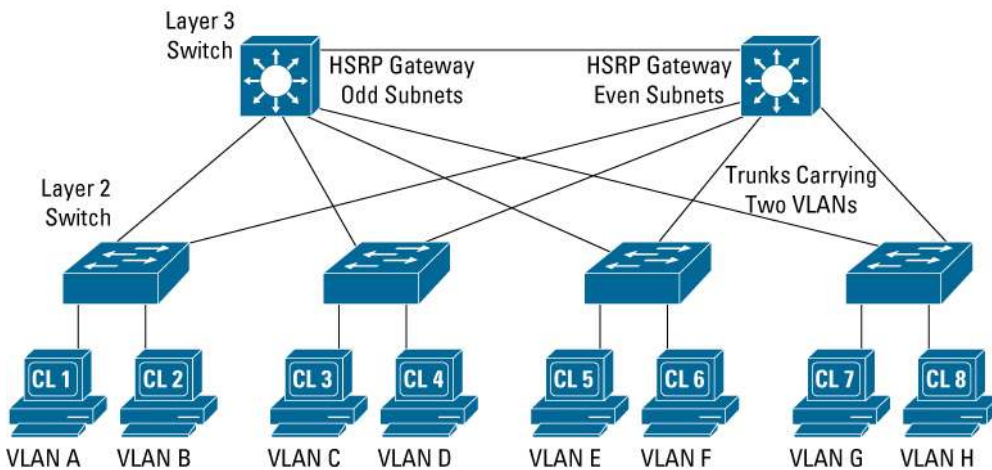


In this configuration, two HSRP groups are configured on the LAN interface of each gateway, which is connected to the Layer 2 switch. Each group is configured such that the router on the right will be active for group 0 and the router on the left will be active for group 1. In this example, off net traffic originating from clients 1 and 3 will go through GW2 while off-net traffic from client 2 goes through GW1. Note the different default gateway assignments for the clients.

However, this configuration carries with it some degree of administrative overhead. The administrative task of dividing clients on a common subnet among multiple default gateways can be tedious, and reduces the plug-n-play characteristic many network administrator desire. Some of this burden is reduced by using smart DHCP servers that help automate the allocation of default gateway IP addresses but even this requires additional effort.

In another variation, we can load-share based on alternating the HSRP-primary function on the redundant layer-3 distribution switches for two or more VLAN/subnets (Figure 5).

Figure 5. Multi-VLAN, Layer 2 Load Balancing Using MHSRP



In this example, each of the local LAN interfaces on the Layer 3 switches are configured as trunks, and each trunk can carry two VLANs. Multiple HSRP groups are defined so that the Layer 3 switch on the left will be active for 'odd' numbered IP subnets while the Layer 3 switch on the right is defined to be active for 'even' numbered IP subnets. The standby HSRP layer-3 switch takes over in the event of failure. Gateway IP addresses for clients are managed on a per-VLAN basis.

The downside with this approach is that in addition to the administrative burden, multiple VLANs and IP subnets must be allocated.

Wants Become Needs

Applications continue to rely increasingly on distributed sources of data and information as they consume more bandwidth. As a result, networking systems continue to evolve with increased interface speeds and packet forwarding rates. The underlying complexity of these systems continues to increase. Companies must focus more on the relationship they have with the business processes, rather than on details of balancing access layer traffic, in order to gain competitive advantage from their networked systems. The network should ideally handle such tasks automatically.

Cisco continues to develop the core technologies that Enterprises depend on in response to customers' wants and needs. GLBP was developed to provide automatic, first-hop gateway load balancing, in order to provide efficient utilization of resources and to reduce administrative costs.

GLBP SOLUTION

The IP protocol requires that endpoints utilize a default gateway to exit a local network and access remote networks. This requirement is the same in both client workstations and data center servers. The endpoint must have prior knowledge of the gateway IP address unless Proxy ARP is used. This is typically that of a router or Layer 3 switch. Endpoints are either statically configured with the IP address of the default gateway or are assigned the address through DHCP upon boot-up. In both cases, the endpoint uses the same default gateway address for all network traffic destined beyond the local network.

In order to forward traffic to the default gateway, the endpoint must perform an IP-ARP resolution to gain the data-link MAC address of the default gateway. The default gateway will respond to the ARP, which effectively informs the endpoint of the gateway's MAC address. The endpoint needs that information in order to forward network traffic to the gateway via a data-link layer transfer.

As Cisco AVVID designs promote resiliency through redundant access layer uplinks and layer-3 gateways, two (or more) possible paths and devices exist for IP gateway services. However, in a standard protocol approach, all endpoints in a common subnet will use the same default gateway and uplink path for all transmissions, which leaves the second possible path unused.

The purpose of GLBP is to *allow load sharing of traffic from access layer endpoints within a common subnet through redundant default gateways, while also providing failure protection*. One way to direct traffic from access layer endpoints to a particular default gateway is to have the gateway respond to the ARP request with its own unique MAC address. Once a particular gateway responds to an ARP request from an endpoint with its own unique MAC address, that endpoint caches the response and will continue to use the discovered gateway for all transmissions destined external to the local subnet. Therefore, if all the redundant Layer 3 gateways in a redundancy group selectively respond to ARP queries in a shared and ordered fashion, workstation traffic exiting the local network will be divided across all possible gateways.

GLBP is an implementation of this general capability. With GLBP, a master controller known as the active virtual gateway (AVG) handles assignment of virtual MAC addresses and responds to ARP requests on behalf of the GLBP redundancy group.

Resiliency services are also provided, so a remaining gateway will assume a failed gateway's endpoint load in addition to its own if there is a failure. The failure remains transparent to the local endpoints. In addition, the protocol supports the ability to 'ignore' the fact that more than one device exists with the capability to accept traffic destined to the same IP address, namely the redundant default gateways.

GLBP builds on the capabilities inherent in HSRP and extends its functionality beyond redundancy, to also include network load sharing among multiple gateways.

SUPPORTED PLATFORMS

GLBP is supported on the following hardware platforms:

- Cisco 1700 Series Routers
- Cisco 2600 Series Routers
- Cisco 3640 Router
- Cisco 3660 Router
- Cisco 3725 Router
- Cisco 3745 Router
- Cisco 7200 Series Routers
- Cisco 7300 Series Routers
- Cisco 7500 Series Routers
- Cisco 7600 Series Routers
- Cisco Catalyst® 6500 Series Switches

Cisco IOS Software is packaged in feature sets that support specific hardware. To get updated information regarding hardware support for this feature, please visit Cisco Feature Navigator at: <http://www.cisco.com/go/fn/>.

This application dynamically updates the list of supported hardware as new hardware support is added for the feature.

GLBP DESIGN CRITERIA

The primary function of the Gateway Load Balancing service is to steer local endpoints to one of many possible default gateways in a load balancing arrangement using IP-ARP mechanisms while using a single virtual IP address for the balanced gateway service. GLBP is the protocol used to implement this service.

The developers of GLBP had a number of specific design goals, which resulted in the following significant implementation achievements:

- **Fast Failover**—GLBP failover occurs immediately after the failure in a gateway is detected. Like HSRP, end stations and applications continue as if no failure had occurred.
- **Low Overhead**—The protocol uses minimal processor and bandwidth resources.
- **Authentication and Security**—The protocol initially implements a very simple authentication scheme. An 8-character string carried in every packet is compared with a configured value, and only packets that match are accepted. MD5 authentication is planned for a future release.
- **Extensibility**—The protocol was designed to be very extensible, due to the packet encoding. This allows for future extensions and enables retroactive compatibility with older releases of Cisco IOS Software.
- **Flexibility**—Several load-balancing algorithms are available for different network configurations and customer requirements.
- **Simple Configuration**—Basic configuration of the protocol is very simple, and is similar to HSRP configuration since many customers are already familiar with it. Each GLBP gateway in a group **MUST** be configured with the same group number and GLBP virtual IP address. All other parameters **MAY** be learned.

GLBP OPERATION

GLBP specifies a protocol that provides load balancing over multiple gateways via a single virtual IP address and multiple virtual MAC addresses.

The members of a GLBP group elect one gateway to be the **Active Virtual Gateway (AVG)** for that group. Other members of the group provide backup for the AVG in case it later becomes unavailable. The AVG assigns a virtual MAC address to each member of the GLBP group. These gateways become the **Active Virtual Forwarder (AVF)** for that virtual MAC address, and assume responsibility for forwarding packets sent to the virtual MAC address.

The AVG is responsible for answering ARP requests for the virtual IP address. Load sharing is achieved by answering these requests with the specific virtual MAC addresses allocated to each of the virtual forwarders.

Addressing

GLBP defines communication among member routers and hosts using real and virtual addresses.

Each GLBP group member must maintain a unique real IP address and a real MAC address. These addresses are used for intra-group communication especially for non-GLBP protocols. For the purposes of a real MAC address, the Burned-in Address (BIA) of the Cisco device interface is used to ensure its uniqueness. The user must assign a unique IP address to each redundant layer-3 entity in the GLBP group.

Each GLBP group member is assigned a unique **virtual MAC address (vMAC)** by the AVG. This MAC is used by local IP endpoints to forward traffic to the specific AVF, which acts as the default gateway. When an IP endpoint ARPs for the default gateway IP address, the unique virtual MAC address will be provided by the AVG in the ARP response. Each IP endpoint is assigned to a particular AVF (and vMAC) according to the defined load balancing method.

Each GLBP group member must be configured with a common **virtual IP address (vIP)**. This address is shared by the group and is used as a default gateway by the endpoints in the local subnet. The virtual IP address is assigned by the network administrator and must be unique. This vIP is

the address that IP endpoints will use as their default gateway. This IP address along with the unique virtual MAC address is installed as a permanent ARP entry in each GLBP group member.

Although multiple gateways will forward traffic destined for the same virtual IP address (as a default gateway), they will be associated with different virtual MAC addresses. The issue of duplicate IP addresses in the subnet will not be a problem provided:

1. The GLBP service is enabled within a Layer 2 switched environment. As ARP responses are directed as Layer 2 unicasts to the requesting station, GLBP group members will not see each other's responses to ARP requests thereby not noticing the duplicate address.
2. The GLBP group members will use the unique real addresses when communicating between each other, but not the virtual IP addresses.
3. The GLBP group members will have permanent ARP entries installed in the ARP cache for the Virtual IP, and their unique virtual MAC addresses to prevent any sort of interference.
4. Any gratuitous ARP mechanisms on the GLBP group members for the virtual IP address are disabled to eliminate conflict between members (failover is an exception that will be addressed).

Multicast and MAC Addresses

GLBP will use the following multicast destination for packets sent to all GLBP group members:

224.0.0.102, UDP port 3222

Virtual MAC addresses will be of the form:

0007.b4yy.yyyy

where yy.yyyy equals the lower 24 bits. These bits consist of 6 zero bits, 10 bits that correspond to the GLBP group number, and 8 bits that correspond to the virtual forwarder number.

0007.b400.0102 @ last 24 bits = 0000 0000 0000 0001 0000 0010 = GLBP group 1, forwarder 2

The protocol allows for up to 1024 GLBP groups and 255 forwarders, but the configuration is currently capped at 4 virtual forwarders per redundancy group. In practice, hardware dependencies relative to the available MAC address filter settings will determine the maximum number of groups and forwarders that can be configured.

AVG Election

The Active Virtual Gateway is determined by GLBP priority. The priority is set using the `glbp <group> priority` interface configuration command. If routers have the same GLBP priority, the gateway with the highest real IP address will become the AVG. HSRP uses an identical operation.

Proxy ARP

Use of GLBP is not recommended in environments where a majority of IP endpoints use proxy ARP. Read on to understand why.

Some IP hosts use proxy Address Resolution Protocol (ARP) to select a router. If a host is running proxy ARP, it sends an ARP request for the IP address of the destination host. A router on the network would reply on behalf of the remote host and would provide its own media access control (MAC) address. With proxy ARP, the host behaves as if the remote host is connected to the same segment of the network.

When proxy ARP is used, a router that sees an ARP request for a remote network will only send an ARP reply if the router has a route to the remote network via a different interface from the one the ARP request was received on. If there is no such route to the network, no ARP reply will be sent.

When GLBP is configured, an additional test occurs before the ARP replay is sent. Therefore, if there is no suitable route it doesn't matter whether GLBP is configured or not.

The ARP reply will be sent with the source MAC address set to the lowest numbered AVF on the interface, if there is a suitable route and a GLBP AVF on the interface that receive the ARP request.

Normally, all GLBP group member routers will have an AVF, so any router that has a suitable route will send an ARP reply. This means there can be multiple ARP replies each with a different source MAC address.

Note: The equivalent case with HSRP is multiple ARP replies with the same source MAC address.

The effect of this is that with proxy ARP enabled, the intended load balancing method is subverted since the IP end point will likely use the MAC address from the last ARP reply received as the destination for off-net packets.

Virtual Gateway Redundancy

Virtual Gateway (VG) Redundancy works the same as HSRP. One gateway is elected the Active VG and another is elected the Standby VG, while remaining gateways are placed in the Listen state. If the Active VG fails, then the Standby VG will assume responsibility for the virtual IP address, and a new Standby VG will be elected from the Listen gateways.

Virtual Forwarder Redundancy

Virtual Forwarder (VF) Redundancy is similar to VG Redundancy in that one of the gateways will takeover traffic forwarding for the virtual MAC address if the active VF fails. However, there are a number of important differences between VF Redundancy and VG Redundancy:

Each gateway in a GLBP group is assigned a virtual MAC address by the AVG (i.e. it becomes a VF for that MAC address). A VF that is assigned a virtual MAC address will become the Active VF for that forwarder instance, and is known as the Primary Virtual Forwarder.

Other gateways in the GLBP group will learn of this VF instance via Hello messages, and will create their own forwarder instances for this forwarder number. These are known as Secondary Virtual Forwarders. Each forwarder runs a timer that is reset to one Hold time whenever a hello message from the active forwarder is received. The active timers running in the other forwarders will expire after one Hold time period if the active forwarder fails. This will cause each of the forwarders in listen state to enter the active state and immediately send out their own active hello message. A lower priority forwarder in active state will return to listen state when it receives a hello from a higher priority active forwarder. Thus, if more than one forwarder becomes active as a result of the failure, all but one of them will return to listen state in the time it takes to send and receive a hello message. The gateway that is the new Active VF is also a secondary VF, so it will be the Primary VF for a different forwarder number. Hosts should be migrated away from a vMAC once a primary virtual forwarder has failed so it can be deprecated unless the PVF comes back online in a reasonably short period of time (a reload, for example). To do this, the Active Secondary VF starts two timers as soon as it changes to the active state—Redirect timer, and Secondary Hold timer. The state of these timers is communicated in the VF Type-Length-Values (TLVs) sent in Hello messages.

The Redirect time is the interval for which the AVG should continue to redirect hosts to the VF MAC address. When this expires, the AVG will stop redirecting hosts to this VF, however the VF continues to forward packets sent to the VF MAC address. The Secondary hold time is the interval that the VF is valid. When this expires then the VF will be removed from all gateways in the GLBP group.

The forwarder number is eligible for re-assignment by the AVG when the AVF is not the Primary Virtual Forwarder.

Timers

There are four main timers that control GLBP operation.

- **Hello time:**

- The approximate period between the Hello messages sent by the GLBP gateway.
- The Hello time is normally learned from the AVG. If the hello time is not learned the manually configured hello time is used. The default value is 3 seconds and the range is 50 milliseconds to 60 seconds.

- **Hold time:**

- Hold time is used to determine if action should be taken to takeover forwarding and/or the AVG function. Each time a hello is received, this timer is re-started. The Hold time should be at least three times the value of the Hello time and must be greater than it.
- The Hold time is normally learned from the AVG. If the Hold time is not learned, the manually configured time is used. The default value is 10 seconds, and the range is from 1 second to 180 seconds.

- **Redirect time:**

- The time for which the Active Virtual Gateway (AVG) continues to redirect hosts to an Active Virtual Forwarder. The objective is to continue handling new ARP requests according to the current load balancing scheme, in anticipation that a failed virtual forwarder will return online. If this occurs within the Redirect time, the returning virtual forwarder will regain his prior load.
- The Redirect time is normally learned from the AVG. If the time is not learned, the manually configured value is used. The default is 5 minutes, and the range is from 1 second to 60 minutes.

- **Secondary Hold Time:**

- The period of time for which a Secondary Virtual Forwarder (SVF) remains valid after the Primary Virtual Forwarder becomes unavailable. An SVF is deleted when the Secondary hold time expires. Once the SVF is deleted, the load balancing algorithm is adjusted to allocate forwarding among the remaining VFs. This time should be longer than the ARP cache age timer of any IP endpoint.
- The Secondary hold time is normally learned from the AVG. If not learned, the manually configured value is used. The default is 1 hour, and the range is from 40 minutes to 18 hours.

The Hello time and Hold time is set using the command:

```
Router(config-if)#glbp group timers [msec] hellotime [msec] holdtime
```

The Redirect time (redirect) and the Secondary Hold time (timeout) is set using the command:

```
Router(config-if)#glbp group timers redirect  
redirect timeout
```

Load Balancing Modes

There will be three types of load balancing methods that can be configured:

- Weighted
- Host dependant
- Round robin

Weighted Load Balancing Algorithm

The amount of load directed to an AVF is dependant upon the weighting value advertised by the gateway containing that AVF.

Each Virtual Forwarder in a gateway uses the current weighting value of that gateway, regardless of how many Virtual Forwarders are active in that gateway.

Host Dependant Load Balancing Algorithm

The host is directed towards a Virtual Forwarder MAC address. In order to determine which Virtual Forwarder MAC address, the MAC address of the host is considered. This ensures that a host will be guaranteed to use the same virtual MAC address as long as that virtual MAC address is participating in the GLBP group. In the current implementation, a host can and will receive a different vMAC in an ARP response if and when the Redirect time expires.

Host dependant load balancing is required in situations where other Cisco IOS Software features use the internal IP Redundancy API. The IP Redundancy API is used for stateful fail over, and this requires each host to be returned the same virtual MAC address each time it ARPs for the virtual IP address. Windows machines, for example, commonly re-ARP every ten minutes. If round-robin load balancing were in place, a Windows machine could switch to using a different gateway every ten minutes. Since stateful failover expects a given IP endpoint to be associated with a specific gateway, this would cause problems. Host-dependent load balancing ensures the same gateway (same vMAC) is used each time it ARPs for the vIP.

Round Robin Load Balancing Algorithm

Each Virtual Forwarder MAC address takes turns being included in address resolution replies for the virtual IP address. Round robin load balancing is **recommended for situations where there are a small number of end hosts**.

If no load-balance algorithm is specified then GLBP will operate in a similar fashion to HSRP, i.e. the AVG will only respond to ARP requests with its own VF MAC address, and all traffic will therefore be directed to the AVG. No load balancing is defined using the following configuration statement:

```
no glbp <glbp-group> load-balancing
```

The load balancing method will be set to default (round-robin) if any load balancing statement is omitted.

GLBP Tracking

GLBP uses Cisco IOS Enhanced Object Tracking, a relatively new software feature. The HSRP interface-tracking feature has been widely deployed and permits the standby router to take over in the event of a failed WAN connection (line-protocol “down”), an event more likely than a router failure. GLBP with Enhanced Object Tracking allows tracking of objects beyond interface tracking.

Cisco IOS Software will deliver Enhanced Object Tracking in phases, so the functionality described here may only be available in certain releases. Table 1 summarized the objects available for tracking, and * indicates that an object that is currently supported. Additional objects for tracking are being considered.

Enhanced Object Tracking is available for HSRP, GLBP, and VRRP.

Table 1. Tracked Objects

Object Type	Object Parameter	Description
Interface (or sub-interface)	Line-protocol*, IP routing*, IPv6 routing, bandwidth, reliability, load	Interface line-protocol is backward compatible with HSRP. Other options extend the function to track the IP routing state, visible bandwidth, reliability, and load.

Object Type	Object Parameter	Description
IP Route	Reachability*, metric value*, metric threshold*	For IP route reachability: If the route exists and the metric is accessible, then the state is “up”. IP routing protocol metrics are scaled into a range from 0 to 255. This scaling factor is configurable. For example, the default resolution for EIGRP is 2560. Therefore, a change in the metric of 2560 will correspond to a change of 1 in the scaled range. The following protocol metrics may be tracked: connected, static, RIP, IGRP, EIGRP, OSPF, and ISIS.
Tracking of Service Assurance Agent Operations	All SA Agent Operations*	The Cisco Service Assurance Agent (SA Agent) is an application-aware synthetic operation agent that monitors network performance by measuring key metrics such as response time, availability, jitter (interpacket delay variance), connect time, throughput, and packet loss. These metrics can be used for troubleshooting, for analysis before problems occur, and for designing future network topologies. The results of SA Agent Operations may be tracked and used to influence GLBP, HSRP, and VRRP behavior.
IPv6 Route	Reachability, metric value	Same as IP route.
Object List	Boolean and*, Boolean or*, Boolean threshold*, Combination of objects/List of objects*	The enhanced tracking process allows for logical OR and logical AND of two or more other tracked objects to create a single new object that can be tracked. The initial tracked objects must have been previously defined.

An additional operation may be specified where ‘up and ‘down’ thresholds are defined. The object state is “up” when the number of “up” objects reaches the ‘up’ threshold, and “down” when the number of “down” objects reaches the ‘down’ threshold.

EXAMPLE CONFIGURATION

A sample configuration is shown in Example 1.

Example 1: Sample GLBP Configuration

Router “Scully”

```
track 30 interface Serial3/0 line-protocol up delay 30
!
interface FastEthernet1/0
ip address 10.44.1.1 255.255.255.0
duplex full
glbp 1 ip 10.44.1.10
glbp 1 weighting 100 lower 95
glbp 1 weighting track 30
glbp 1 forwarder preempt delay minimum 0
```

Router “Lonegunman”

```
track 30 interface Serial3/0 line-protocol up delay 30
!
interface FastEthernet1/0
 ip address 10.44.1.2 255.255.255.0
 duplex full
 glbp 1 ip 10.44.1.10
 glbp 1 priority 95
 glbp 1 weighting 100 lower 95
 glbp 1 weighting track 30
 glbp 1 forwarder preempt delay minimum 0
```

In the example configuration, GLBP is enabled on both routers on interface Fast Ethernet 1/0. The virtual IP address for GLBP group 1 is 10.44.1.10. Downstream IP hosts will point to this address as their default gateway.

Note that the GLBP priority for router “Lonegunman” was set to ‘95’. The default priority is 100. This change will cause router “Scully” to become the AVG since its priority is higher.

GLBP is configured in both routers to track object number 30, which, in turn, defines the object as the line-protocol for serial interface 3/0. If the line-protocol for S3/0 goes ‘down’, the GLBP weighting will be decremented by ‘10’. This is the default amount to decrement and can be modified with the **glbp 1 weighting track 30 decrement** command. When the weighting crosses the lower threshold, in this case ‘95’, GLBP will change the state of the virtual forwarder and the standby VF will take over.

When the line-protocol for S3/0 is restored, GLBP will delay establishing an AVF for 30 seconds as specified by the **up delay** on the **track** statement.

Please see the Cisco IOS Software documentation for more detailed information on the GLBP commands.

Troubleshooting Tools

Cisco IOS Software contains show and **debug** commands to allow network operators to determine the status of GLBP. These command are useful in verifying the operation of GLBP and that the desired design goals are being met relative to tracked objects and forwarder states. Some sample show and debugs are depicted in Example 2 and Example 3.

Example 2: GLBP Show Commands

Scully	Lonegunman
scully#sh glbp	lonegunman#sh glbp
FastEthernet1/0 - Group 1 ∟ 1.1	FastEthernet1/0 - Group 1 ∟ 1.2
State is Active	State is Standby
2 state changes, last state change 5d20h	13 state changes, last state change 5d20h
Virtual IP address is 10.44.1.10 ∟ 2.1	Virtual IP address is 10.44.1.10 ∟ 2.2
Hello time 3 sec, hold time 10 sec	Hello time 3 sec, hold time 10 sec
Next hello sent in 0.252 secs	Next hello sent in 1.464 secs
Redirect time 600 sec, forwarder time-out 3600 sec	Redirect time 600 sec, forwarder time-out 3600 sec
Preemption disabled	Preemption disabled
Active is local ∟ 3.1	Active is 10.44.1.1, priority 100 (expires in 7.196 sec) ∟ 3.2
Standby is 10.44.1.2, priority 95 (expires in 8.524 sec)	Standby is local
Priority 100 (default) ∟ 4.1	Priority 95 (configured) ∟ 4.2
Weighting 100 (default 100), thresholds: lower 1, upper 100	Weighting 100 (default 100), thresholds: lower 1, upper 100
Load balancing: round-robin ∟ 5.1	Load balancing: round-robin ∟ 5.2
There are 2 forwarders (1 active)	There are 2 forwarders (1 active)
Forwarder 1 ∟ 6.1	Forwarder 1 ∟ 6.2
State is Listen	State is Active
2 state changes, last state change 5d20h	7 state changes, last state change 5d20h
MAC address is 0007.b400.0101 (learnt) ∟ 7.1	MAC address is 0007.b400.0101 (default) ∟ 7.2
Owner ID is 0030.969a.8c1c	Owner ID is 0030.969a.8c1c
Redirection enabled, 598.524 sec remaining (maximum 600 sec)	Preemption enabled, min delay 30 sec
Time to live: 3598.524 sec (maximum 3600 sec)	Active is local, weighting 100 ∟ 8.2
Preemption enabled, min delay 30 sec	Forwarder 2 ∟ 9.2
Active is 10.44.1.2 (primary), weighting 100 (expires in 8.524 sec)	State is Listen
Arp replies sent: 429	2 state changes, last state change 5d20h
Forwarder 2 ∟ 9.1	MAC address is 0007.b400.0102 (learnt)
State is Active	Owner ID is 0002.4a29.e01c
1 state change, last state change 5d20h	Time to live: 3598.992 sec (maximum 3600 sec)
MAC address is 0007.b400.0102 (default)	Preemption enabled, min delay 30 sec
Owner ID is 0002.4a29.e01c	Active is 10.44.1.1 (primary), weighting 100 (expires in 8.992 sec)
Redirection enabled	
Preemption enabled, min delay 30 sec	
Active is local, weighting 100	
Arp replies sent: 429	

The sample **show glbp** commands in Example 2 are from a pair of Cisco 7200 Series Routers, each configured for GLBP on interface fa1/0. The display on the left is from a router called “scully” and the display on the right is from a router called “lonegunman”.

As the display demonstrates, quite a bit of information can be obtained from the show command. The statuses shown in the displays are from the perspective of the router where the command is entered. The notations in the example are described below:

∟ 1: The display shows the interface(s) that have GLBP configured and lists the defined GLBP groups. The next line indicates the current GLBP group state. Active means the router is acting as the AVG for this GLBP group. So at 1.1 we see that “scully” is the AVG and at 1.2 we see that “lonegunman” is the Standby. “Lonegunman” will become the AVG if “scully” fails.

∟ 2: The common virtual IP address used for the GLBP redundancy group.

∟ 3: An indication that the AVG is local or if this router is acting as the Standby, as well as which other router is performing the AVG or Standby function for the GLBP group.

∟ 4: An indication of the current priority set for the gateway.

∟ 5: Shows the load balancing method being used.

∟ 6: Display for each of the active forwarders in the GLBP group. There are two active forwarders, 1 and 2. For “scully” forwarder 1.1 (group.forwarder notation) is in “listen” state and forwarder 1.2 is “active”.

∟ 7: Shows the virtual MAC address and, on the next line, the real MAC address associated with the assigned owner (primary virtual forwarder) of this vMAC. The owner and vMAC was previously assigned by the AVG.

∟ 8: Shows the current weighting value for this VF.

∟ 9: Similar display for the next virtual forwarder.

Example 3: Sample Debug Command

```
scully#debug glbp event
GLBP Events debugging is on
scully#
Apr 19 11:50:54.639: GLBP: Fa0/0 API active virtual address 172.26.64.1 not found
Apr 19 11:50:57.363: GLBP: Fa0/0 API active virtual address 172.26.64.1 not found
Apr 19 11:51:20.295: GLBP: Fa1/0 1 Track 30 object changed, state Up -> Down ∟ 1
Apr 19 11:51:20.295: GLBP: Fa1/0 1 Weighting 100 -> 90
Apr 19 11:51:21.563: %LINK-3-UPDOWN: Interface Serial3/0, changed state to down ∟ 2
Apr 19 11:51:22.131: GLBP: Fa1/0 1.2 Active: i/Hello rcvd from higher pri Active router (135/10.44.1.2) ∟ 3
Apr 19 11:51:22.131: GLBP: Fa1/0 1.2 Active -> Listen ∟ 4
Apr 19 11:51:22.131: %GLBP-6-FWDSTATECHANGE: FastEthernet1/0 Grp 1 Fwd 2 state Active -> Listen ∟ 5
Apr 19 11:51:22.135: GLBP: Fa1/0 API MAC address update ∟ 6
Apr 19 11:51:22.563: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial3/0, changed state to down
Apr 19 11:51:42.939: GLBP: Fa0/0 API active virtual address 172.26.64.1 not found
Apr 19 11:51:57.891: GLBP: Fa0/0 API active virtual address 172.26.64.1 not found
Apr 19 11:52:03.495: %LINK-3-UPDOWN: Interface Serial3/0, changed state to up
Apr 19 11:52:04.495: %LINEPROTO-5-UPDOWN: Line protocol on Interface Serial3/0, changed state to up
Apr 19 11:52:34.295: GLBP: Fa1/0 1 Track 30 object changed, state Down -> Up ∟ 7
Apr 19 11:52:34.295: GLBP: Fa1/0 1 Weighting 90 -> 100 ∟ 8
Apr 19 11:52:35.127: GLBP: Fa1/0 1.2 Listen: k/Hello rcvd from lower pri Active router (135/10.44.1.2)
Apr 19 11:52:35.127: GLBP: Fa1/0 1.2 Listen -> Active
Apr 19 11:52:35.127: %GLBP-6-FWDSTATECHANGE: FastEthernet1/0 Grp 1 Fwd 2 state Listen -> Active
```

In Example 3, the output from **debug glbp event** command is shown. While the debug was active, the serial link from router “scully” was failed by shutting the remote end.

∟ 1: An indication that the tracked object (serial 3/0) was observed to have changed state. The next line shows the GLBP group 1 weighting being from 100 decremented by 10 to 90.

∟ 2: This is the router log message indicating that Serial 3/0 has gone “down”.

∟ 3: The message that appears at 11:51:22.131 indicates that GLBP received a “Hello” message from another GLBP group member. This is due to the fact that this router’s priority has been reduced according to the weighting track decrement amount. This is an indication that another router can take over forwarding.

∟ 4: The state change for this router’s virtual forwarder 1.2 from “active” to “listen”.

∟ 5: This is a log message written by GLBP indicating the forwarder state change.

∟ 6: The “API” messages refer to internal Cisco IOS Software process exchange events.

∟ 7: An indication that the tracked object has come back “up”. Note the time is 11:52:34:295. This occurs about 30 seconds after the line protocol actually came back up for the object. This delay was specified using the **track <object no.> up delay 30** command.

∟ 8: Next the weight change can be seen and “scully” again becomes the AVF for the vMAC.

Other debug options are possible including errors, events, and packets. See the Cisco IOS Software documentation for a full set of debug commands and syntax.


DEPLOYMENT SCENARIOS FOR GLBP

The following scenarios will benefit from the GLBP service as previously described. There will be some scenarios in which it is inappropriate to deploy GLBP or interaction with other IP services and protocols will render impractical GLBP (in its current form).

Small Remote Site with Redundant Routers

GLBP is ideal for remote sites that have high availability requirements for access to applications and servers at a central location (see Figure 6). By adding GLBP to the edge routers, upstream bandwidth can be more efficiently utilized, improving performance and productivity during non-fault conditions. Fault protection from router and link or circuit problems is guaranteed by simultaneous use of redundant facilities whereby each path can takeover the load if a failure occurs.

Figure 6. Small Remote Site




Variation Using Redundant Service Providers

It is possible to vary this design for a remote site by using two different service provider networks for WAN connectivity. This increases redundancy benefits and protects against problems specific to a particular provider over and above simple router and access link failures. Special considerations apply in environments where network address translation (NAT) is used, as the ARP-based load sharing of GLBP may redirect client traffic to the alternate path. Currently, the use of NAT and more specifically, Stateful NAT is not recommended with GLBP.

Campus Access Layer Design

GLBP can be employed at the campus access layer to make better use of upstream Fast Ethernet or Gigabit Ethernet links without the need to manage and segment clients among different layer-3 default gateway addresses and VLANs (see Figure 7). GLBP deployment offers the benefits of redundancy and load sharing, as well as reduced dependency on STP and multiple VLANs per access switch.

Figure 7. Campus Design with GLBP



Asymmetric Routing

Asymmetric routing occurs when packets are sent from a source to destination over one path while return traffic follows a different path. This will often be the case with GLBP, since traffic is intentionally being shared over multiple upstream paths. Asymmetric routing can cause excessive flooding of unicast IP packets. This adverse condition results from the MAC address of downstream hosts being aged out of the switch CAM (typical default time is 5 minutes). The CAM aging time can be increased to a higher value to avoid this condition. One suggestion for the Cisco Catalyst 6500 is to make the CAM aging time equal to the ARP timeout for the MSFC. For example, set the Layer 2 CAM entry aging time in distribution layer switches to the same duration as the ARP timeout using the command:

```
set cam agingtime 1-1000 14400
```

This will set the time to 4 hours, same as default ARP cache timeout for MSFC and will minimize any flooding of IP unicast traffic when packets are never received for a given MAC.

Cisco Catalyst 6500 Series MAC Address Details

GLBP “reserves” four hardware MAC filter entries. This limits the number of virtual forwarders to four. This is a somewhat arbitrary cap but should be sufficient for customers. Cisco has not received requests for support of additional active forwarders.

There is an additional restriction for the Supervisor II with the MSFC2/PFC2. This hardware limits the number of GLBP groups that may be defined to one (i.e. only a single GLBP group ID can be used per system). This is due to the MAC lookup method and the GLBP MAC address format. However, this single group ID may be reused on all VLANs.

There is a 1024 group limit for the Sup720. Again, this limit is arbitrary and the protocol design actually allows for 4096.

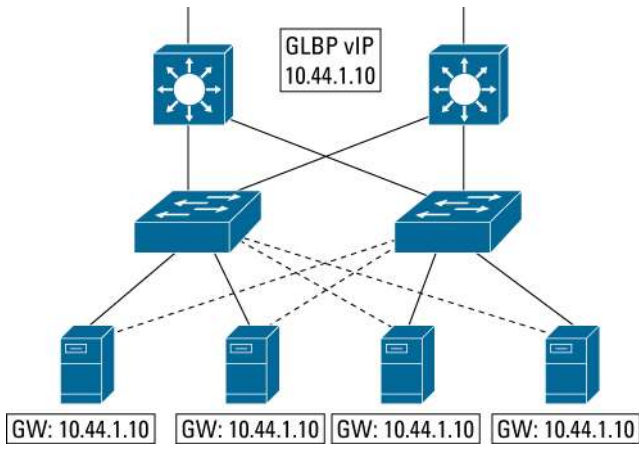
Table 2. C6500/7600 MAC Restrictions

Cisco IOS Software Release	Switching Product	Group/Forwarder Limit
12.2(17d)SXA and later	C6500/7600 Sup720 MSFC3	1024/4
12.2(17d)SXB and later	C6500/7600 Supervisor II MSFC2	1/4

Server Farm Design

GLBP can be used to share outbound network traffic from simple server farms. GLBP is best used in environments that do not require additional server load balancing features and function such as that provided by Cisco IOS Server Load Balancing (IOS-SLB), Cisco Content Switching Module (CSM), and the Cisco Content Switching Services (CSS) family of products (see Figure 8).

Figure 8. Simple Server Farm Design with GLBP



GLBP BENEFIT SUMMARY

Gateway Load Balancing Protocol allows enterprises to get more benefit from their network resources, links and equipment, when there is no failure, while still providing for the desired redundancy for high availability. The benefit comes in terms of network performance and resource efficiency.

Reduced Queuing

By using multiple available paths upstream from the routers or Layer 3 switches running GLBP, output queues may be reduced. With HSRP or VRRP, only a single path is used while others stand by idle unless multiple groups and gateways are configured. The single path may encounter higher output queue rates during peak times leading to lower performance from higher jitter rates. The impact of jitter is reduced and overall performance is increased, because more upstream bandwidth is available and additional output queues are used.


Resource Efficiency

The network exists to service the enterprise, while offering strategic advantage and higher productivity. The enterprise depends on its network to be highly available and to provide high performance. The resources employed to achieve these goals come at significant cost. Therefore, efficient use of these resources to provide direct benefit in terms of productivity and return on investment is important to all organizations.

GLBP can increase productivity by allowing more efficient use of resources. The ability to utilize multiple upstream paths when available has a direct impact on the return the enterprise sees from the investment they've made in network resources. By increasing usable bandwidth, resources are used more efficiently. With higher available bandwidth, file transfer times are lowered, leading to increased productivity.

As an illustration of improved resource efficiency and performance, consider the example shown in Figure 9. A small site with two routers and Layer 2 switches can be seen here. Each upstream link was clocked by the connecting router at 1,612,800 bps. As a comparison, each router was configured for HSRP on fast Ethernet interface 2/0 and with GLBP on fast Ethernet 1/0.

Figure 9. GLBP Deployment, Small Remote Site



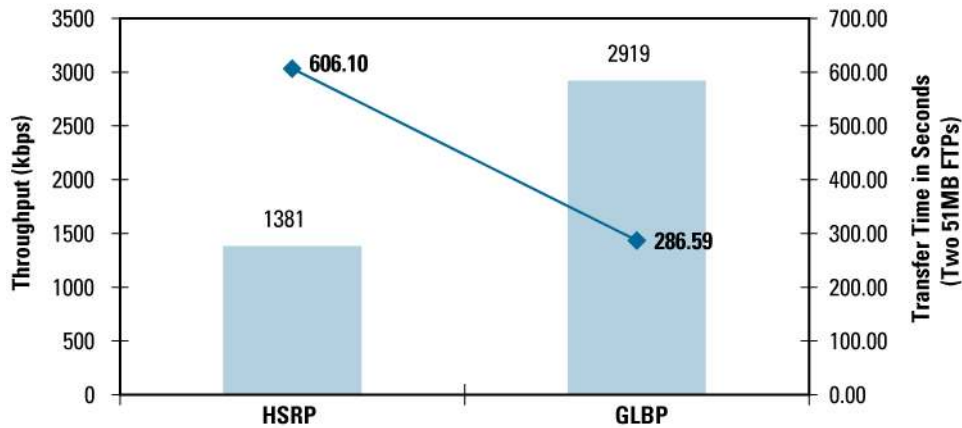
The PCs connected to the switch labeled “oak” were each configured with a default gateway corresponding to the GLBP virtual IP address, 10.44.1.10. Similarly, the PCs connected to the switch labeled “pine” were configured with a default gateway corresponding to the HSRP standby address 10.44.2.10.

Due to the nature of HSRP, the PCs on “pine” will only be able to utilize a single upstream link, the “active” link, while the remaining link remains idle in anticipation of failure to the primary path.

The PCs on “oak”, on the other hand, will have both paths available and will use the path assigned by the AVG according to the GLBP load sharing mechanism defined, in this case, round-robin.

To illustrate the benefit, two concurrent files transfers were conducted from each of the PCs across the WAN to a server, first the ones on switch “pine” and then the ones on switch “oak”. As the graph in Figure 10 shows, throughput increased in proportion to the additional available bandwidth when using GLBP while file transfer times decreased.

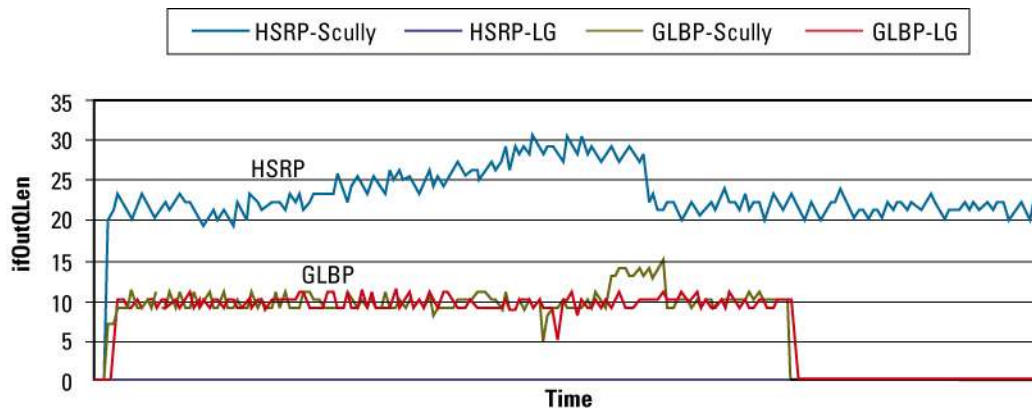
Figure 10. Gateway Redundancy Comparison



In this simple test of two simultaneous transfers of an approximately 51MB file using FTP, throughput increased from 1381 kbps to 2,919 kbps and the average time to transfer the file decreased from 606.1 seconds to 286.6 seconds.

Network performance relative to the size of the output queue observed at the WAN interface also improved when GLBP was compared to HSRP. During the file transfer test, the output queue length was monitored at 2-second intervals. The result is shown in the graph in Figure 11.

Figure 11. Output Queue Comparison



The graph shows that when HSRP is in use, only the single upstream path is utilized so the output queue length for the WAN interface on router “scully” can be seen to range from 20 to 30 packets during the test. In contrast, when GLBP is in use, two paths are utilized and the output queue lengths for the WAN interfaces on each upstream router, “scully” and “lonegunman” or “LG” both ranged much lower.

Again, this illustrates the performance improvement, both in terms of efficiency and performance that GLBP has over HSRP when used certain environments.

SUMMARY

Cisco IOS Software contains many features specifically designed to make the network more resilient and, therefore, increase the availability of access to the applications important to your business. GLBP provides first-hop, Layer 3 redundancy service while sharing the network load over multiple devices and upstream paths. The result is better performance, lower administrative costs, and higher return on your network investment.

DEFINITIONS

Address Resolution Requests/Replies

IPv4 ARP request/reply or IPv6 neighbor solicitation/advertisement.

GLBP Group

A GLBP group consists of one or more GLBP gateways configured with the same GLBP group number.

GLBP Gateway

A gateway or router running the Gateway Load Balancing Protocol. It may participate in one or more GLBP groups.

Virtual IP Address (vIP)

An IPv4 address or IPv6 prefix. There **MUST** be only one virtual IP address configured for each GLBP group. The virtual IP address **MUST** be configured in at least one member of the GLBP group. Other members **MAY** learn the virtual IP address from Hello messages.

Virtual MAC Address

A MAC address that a host may receive when it issues an address resolution request for the virtual IP address. There **MAY** be multiple virtual MAC address for each GLBP group.

Virtual Gateway (VG)

An abstract entity within a GLBP gateway that may assume responsibility for operation of the protocol.

Active Virtual Gateway (AVG)

One Virtual Gateway within a particular GLBP group is elected the Active Virtual Gateway (AVG), and is responsible for operation of the protocol.

Virtual Forwarder (VF)

An abstract entity within a GLBP gateway that may assume responsibility for a virtual MAC address.

Active Virtual Forwarder (AVF)

One Virtual Forwarder within a particular GLBP group is elected the Active Virtual Forwarder (AVF), and is responsible for forwarding packets sent to a particular virtual MAC address. There may be multiple Active Virtual Forwarders in a GLBP group.

Primary Virtual Forwarder (PVF)

A Virtual Forwarder that has been assigned the virtual MAC address by the Active Virtual Gateway. When available, a Primary Virtual Forwarder will always become the Active Virtual Forwarder for a virtual MAC address.

Secondary Virtual Forwarder (SVF)

A Virtual Forwarder that has learned the virtual MAC address from a Hello message.

REFERENCES

Cisco IOS IP Services

<http://www.cisco.com/warp/public/732/Tech/ipservices/>

Cisco IOS High Availability

<http://www.cisco.com/go/availability/>



Corporate Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-4000
800 553-NETS (6387)
Fax: 408 526-4100

European Headquarters

Cisco Systems International BV
Haarlerbergpark
Haarlerbergweg 13-19
1101 CH Amsterdam
The Netherlands
www-europe.cisco.com
Tel: 31 0 20 357 1000
Fax: 31 0 20 357 1100

Americas Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA
www.cisco.com
Tel: 408 526-7660
Fax: 408 527-0883

Asia Pacific Headquarters

Cisco Systems, Inc.
168 Robinson Road
#28-01 Capital Tower
Singapore 068912
www.cisco.com
Tel: +65 6317 7777
Fax: +65 6317 7799

Cisco Systems has more than 200 offices in the following countries and regions. Addresses, phone numbers, and fax numbers are listed on the **Cisco Web site at www.cisco.com/go/offices.**

Argentina • Australia • Austria • Belgium • Brazil • Bulgaria • Canada • Chile • China PRC • Colombia • Costa Rica • Croatia • Cyprus • Czech Republic • Denmark • Dubai, UAE • Finland • France • Germany • Greece • Hong Kong SAR • Hungary • India • Indonesia • Ireland • Israel • Italy • Japan • Korea • Luxembourg • Malaysia • Mexico • The Netherlands • New Zealand • Norway • Peru • Philippines • Poland • Portugal • Puerto Rico • Romania • Russia • Saudi Arabia • Scotland • Singapore • Slovakia • Slovenia • South Africa • Spain • Sweden • Switzerland • Taiwan • Thailand • Turkey • Ukraine • United Kingdom • United States • Venezuela • Vietnam • Zimbabwe

Copyright © 2004 Cisco Systems, Inc. All rights reserved. CCIP, CCSP, the Cisco *Powered* Network mark, Cisco Unity, Follow Me Browsing, FormShare, and StackWise are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn, and iQuick Study are service marks of Cisco Systems, Inc.; and Aironet, ASIST, BPX, Catalyst, CCDA, CCDP, CCIE, CCNA, CCNP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, the Cisco IOS logo, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Empowering the Internet Generation, Enterprise/Solver, EtherChannel, EtherSwitch, Fast Step, GigaStack, Internet Quotient, IOS, IP/TV, iQ Expertise, the iQ logo, iQ Net Readiness Scorecard, LightStream, Linksys, MGX, MICA, the Networkers logo, Networking Academy, Network Registrar, *Packet*, PIX, Post-Routing, Pre-Routing, RateMUX, Registrar, ScriptShare, SlideCast, SMARTnet, StrataView Plus, Stratm, SwitchProbe, TeleRouter, The Fastest Way to Increase Your Internet Quotient, TransPath, and VCO are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Web site are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (0402R) 204025_ETMG_SH_09.04

