



W H I T E P A P E R

Enhancing Web User Experience with Global Server Load Balancing

- **GSLB OPERATION OVERVIEW**
- **MAJOR COMPONENTS**
- **DISTRIBUTED SITE MONITORING**
- **DISTRIBUTED SITE STATE PROTOCOL**
- **INTERNET TOPOLOGY AWARENESS**
- **DNS AUTHORITATIVE NAME SERVER**
- **ALTEON WEBSYSTEMS GSLB ADVANTAGES**

Alteon WebSystems, Inc.

50 Great Oaks Boulevard
San Jose, California 95119
408-360-5500
408-360-5501 fax

<http://www.alteon.com>

W H I T E P A P E R

Global server load balancing (GSLB) allows web hosters, portals and enterprises to distribute content and services geographically. Dispersing content and services offers a number of advantages, including:

- Allowing users to be automatically directed to content from servers located in their own geographic region thus reducing response times and decreasing the use of expensive international data connections.
- Directing users away from congested networks and servers, enhancing the users' experience.
- Increasing fault-tolerance and availability by allowing multi-site content and service deployment, guarding against failures in the event of local or regional network outages, power outages or natural disasters.

Alteon WebSystems Web switches running GSLB direct user requests to the "best site" to service the requests using three criteria:

- 1) site health,
- 2) site proximity and
- 3) response time required to retrieve specified content.

Global server load balancing is explained here in the context of HTTP and the World Wide Web but is by no means limited to HTTP. Any service that can be load balanced with Alteon WebSystems' Web switches can operate with GSLB.

It's important that the reader be familiar with the concepts of Alteon WebSystems' local server load balancing operation detailed in our white paper "Next Steps in Server Load Balancing." This paper is available at http://www.alteonwebsystems.com/products/white_papers/SLB/index.shtml.

GSLB OPERATION OVERVIEW

When client Z loads their browser and enters the URL: <http://www.Alteon.com> (see Figure 1), the system sends a DNS getByHostname query to the client's local DNS server, asking for the IP address that represents www.Alteon.com.

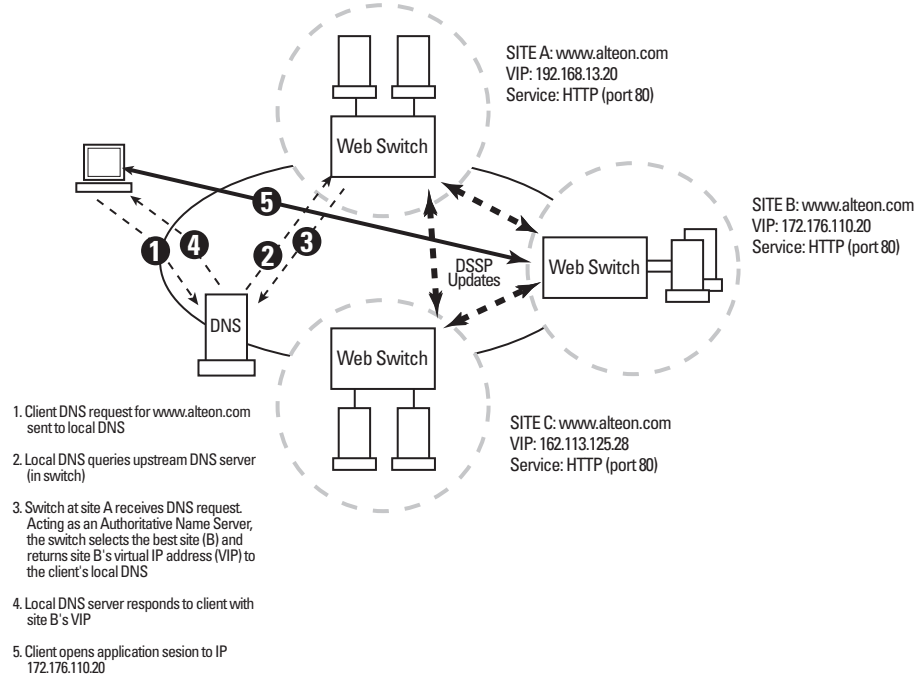
The local DNS server examines its DNS cache to determine if it already knows about this particular domain name and host. If it doesn't, the local DNS server hands the request off to the appropriate upstream DNS server.

The DNS query is either responded to by an upstream DNS server's cache or is passed on until the request arrives at a DNS server embedded in one of the Web switches at site A, B, or C (in this case site A). Which site ultimately receives the request is determined by a myriad of DNS configuration parameters.

The Web switches at site A, B and C are configured to be "distributed sites" and all can act as Authoritative Name Servers for the domain www.Alteon.com. Each can respond directly to DNS queries with IP addresses that represent that domain.

For example, if site A receives the DNS query, the IP address that site A returns represents a Virtual IP (VIP) address for one of the sites hosting the requested content (in this case site B).

FIGURE ONE — GSLB Operation



In the example above, assume that site A returns the IP address 172.176.110.20. The client receives the DNS query response from its local DNS server indicating that 172.176.110.20 is the IP address for www.Alteon.com. It then opens a TCP Port 80 connection to 172.176.110.20, the VIP address running at site B. Now, the client is communicating with the Alteon WebSystems Web site with content from site B.

So, how did site A determine that site B was the right site to handle the client's request? How is site B "better" than the other two possible sites, including the one that responded to the DNS query? With GSLB, three criteria are used to determine to which site DNS will direct the client:

- Site health
- Geographic location of the client and sites(s)
- Measured site response time

GSLB develops an ordered list of sites that DNS uses when responding to client requests. The above criteria are used to determine if and where on the list a site appears (this is detailed later).

What happens if the site to which the client has been pointed suddenly experiences a failure or is overloaded? Assuming the Web switch running GSLB and its Internet connection are up, the Web switch issues an HTTP Redirect back to the client, telling it to go to a different site.

This occurs when a VIP no longer has any healthy real IP addresses (RIPs) or when an HTTP request is sent to real servers that have reached their respective maximum connection thresholds.

MAJOR COMPONENTS

GSLB consists of four major components that run on each Web switch in the GSLB group:

- *Distributed Site Monitoring*—where a Web switch at each site performs Layer 4 health checking (with content verification as an option) on all other peer remote sites. This determines the health and response time of servers and applications at each site.
- *The Distributed Site State Protocol (DSSP)*—used to exchange health, load, response time and throughput information between sites through both periodic updates during normal operation and triggered updates when a significant event occurs.
- *Internet Topology Awareness*—where a Web switch acting as an Authoritative Name Server examines DNS requests and considers geography when responding.
- *A DNS Authoritative Name Server*—responds to DNS requests directed to that site.

DISTRIBUTED SITE MONITORING

A Web switch at each distributed site performs periodic health and response time checks of each defined Remote Real IP (RIP) addresses. These remote RIPs (i.e. devices participating in the GSLB operation) typically correspond to VIP addresses running in Web switches at peer sites being load balanced by GSLB. By executing configurable, iterative health checks to each remote RIP, a site learns about its peer sites' server, application and content availability and response time.

Each health check consists of open and closing a TCP connection for a configured application. For application/protocols where the Web switch supports content health checking (HTTP, FTP, NNTP, DNS, SMTP and POP3), content access can also be configured as part of the health check. Content is accessed based on the content configuration (URL, filename, etc.) defined by the system administrator.

When content-based health checking is used, response time is defined as the time from when the Web switch issues the request to open the connection to the time it closes the connection, including the time needed to retrieve the content. Without content-based health checking, the time needed to retrieve content is not a factor.

Each Web switch performs this health and response time check for each defined remote RIP, each corresponding to a VIP address running in a Web switch at another site.

For instance, if site A sees four other sites and there are five VIPs defined on the Web switch at site A, (each having corresponding remote RIPs at each site), then the Web switch at site A performs 20 health and response time checks during the health check interval (4 sites times 5 remote RIPs at each site).

It's important to note that remote RIP health checks don't stop at the Web switch hosting the remote VIP. The remote Web switch passes them through to a server or servers behind the Web switch.

If the Web switch is in front of a group of load balanced servers, the health checks are distributed across the servers in accordance with the configured load balancing metric. As a result, remote health checks determine the

availability of not only remote Web switches but also the server, applications and if configured, the content behind the Web switches.

If a Web switch flags a remote RIP as down because it does not respond to health checks, the Web switch:

- No longer considers the site eligible for connection handoffs and stops using the remote Web switches VIP address as a target for DNS responses.
- Notifies all other distributed sites that the site is not responsive. Each distributed site may then test to see if the site is responsive and act accordingly.

DISTRIBUTED SITE STATE PROTOCOL

The Distributed Site State Protocol (DSSP) is a light-weight protocol used to communicate health and response time information from one distributed site to every other distributed site. Each DSSP packet communicates:

- Response times for each peer site as measured by the site transmitting the DSSP packet.
- Remaining site capacity (connections available per VIP address) of the transmitting site.
- Status of the transmitting site.

Each site uses the information communicated by the DSSP, plus its own response time checking results, to construct a table of response times for all sites as measured by all sites. This information is, in turn, used to calculate the desired relative traffic distribution between the distributed sites, including itself. For example, the sites might determine that:

- Site A should receive 20% of all traffic.
- Site B should receive 10% of all traffic.
- Site C should receive 10% of all traffic.
- Site D should receive 20% of all traffic.
- Site E should receive 10% of all traffic.
- Site F should receive 30% of all traffic.

The DNS authoritative name server in each Web switch uses these percentages to determine how often each site's VIP address should be included in responses it sends to downstream DNS servers.

The advantages to this algorithm include:

- Sites that perform the best will generally receive more connections than other sites, but not all of the connections. This prevents traffic spikes from overloading individual sites.
- The traffic will be averaged across the top sites, providing consistently good response times and user experiences.

- The sites that are seen as poorly performing by all other sites (an indication of a real problem) will tend to receive few or no connections, providing relief while they process their existing load or corrective action is performed.
- If every site is performing well (including WAN links, servers, etc.) then it's likely that each site will receive an equal distribution of traffic over time. This ensures that sites don't get overloaded but also perform their share of the work.

In addition to regular updates, Web switches send DSSP triggered updates under the following exception conditions:

- The Web switch is no longer able to communicate with a remote RIP.
- The Web switch experiences a local resource constraint, such as all servers have reached their maximum connections limit or no real servers available for a VIP.

DSSP triggered update contains all of the information in a regular update.

INTERNET TOPOLOGY AWARENESS

In addition to site health and response time, GSLB takes geographic information into account when determining which distributed site should handle a request—for instance, if there are five sites that host content for a given host and domain name, one each in San Jose (West-U.S.); Atlanta (East-U.S.); Ecuador (South America); Paris, France; and Tokyo, Japan.

In general, associating users within a geographic domain (country, continent) with servers in that domain optimizes the user's experience (unless the "nearby" site is down or overloaded).

With this in mind, users in Europe will generally be served by the Paris site while users in Chile will be served by the Ecuador site. Having a user in Japan come to the Atlanta site for content would waste expensive international bandwidth and cause unnecessary response delays to the user. Switches at distributed sites also consider geography when responding to DNS requests.

When a Web switch receives a DNS request, it recognizes the geographic source of the request by inspecting the Source IP address of the request. It then consults the relative traffic distribution table (described later) for that geographic area to determine which site within the area the DNS response should indicate.

For example, if the requesting host is located somewhere in the Pacific Rim area, it will be pointed to the server in Tokyo. If the requesting host is located somewhere in the United States, the Web switch will consult the relative traffic distribution traffic table for the U.S. to determine if the host should be pointed to Atlanta or San Jose.

DNS AUTHORITATIVE NAME SERVER

Ultimately, GSLB is accomplished by the DNS Authoritative Name Server running in the Web switches at distributed sites returning the appropriate IP address to downstream DNS servers.

For example, when a client enters a URL into their browser for a particular hostname (represented by several VIPs scattered throughout the U.S.), their system sends a DNS `getByHostname` query to their local DNS server, asking for the IP address representing that domain name and host.

The local DNS server then examines its DNS cache to determine if it already knows about this particular domain name and host. If it doesn't know about the hostname, it hands off the request to the next appropriate DNS server. The DNS query is either responded to from that DNS server's cache or is passed on until the request comes to a Web switch at a distributed site running GSLB.

When a Web switch at a distributed site receives a DNS query to resolve a hostname from a downstream DNS server, it determines to which geographic region the requesting host belongs. It then checks to see if any healthy distributed sites are present in that region. If there are none, it looks for healthy distributed sites in other regions.

Otherwise, the Web switch provides a DNS response containing an IP address based on the relative traffic distribution traffic table for that region. The IP address will change from response to response based on the percentages in the relative distribution traffic table.

ALTEON WEBSYSTEMS GSLB ADVANTAGES

While competitive solutions for distributing load across geographically distributed servers exist, none offer the full range of advantages supported by Alteon Websystems' GSLB. These advantages include:

- 1) Local server load balancing, global server load balancing, application redirection, Layer 2 and 3 switching within a single platform. This enables additional applications such as Web cache redirection, DNS redirection, firewall load balancing and router load balancing. Today, no competitive product can match this level of integration and flexibility.
- 2) GSLB directs users to the best performing sites within a geographic region. Competitive solutions that rely on metrics such as router hops fail to consider important factors such as networks congestion and server load when making their load balancing decisions. Only Alteon WebSystems' GSLB effectively takes these factors into consideration.
- 3) Intelligent load distribution funnels most traffic to the best performing sites without overwhelming them.
All resources are effectively used, improving the user's experience.
- 4) Users are automatically redirected to the next best site when all servers at a site are down or congested, improving service and content availability.

SUMMARY

Global Server Load Balancing (GSLB) allows Web hosters, portals and enterprises to enhance users' Web experiences by reducing response times and increasing service and application availability.

At the same time, organizations deploying GSLB benefit through the decreased use of expensive international data connections. GSLB is an optional software capability on Alteon Web switches such as the ACEdirector 2, ACEdirector 3 and ACEswitch 180. GSLB running on these Web switches directs user requests to the “best site” to service the requests. The “best site” is determined by monitoring site health, site proximity and response time.