APPLICATION NOTE

GLOBAL SERVER LOAD BALANCING WITH SERVERIRON

Growing Global

Simply by connecting to the Internet, local businesses transform themselves into global "eBusiness" enterprises that span the world. When the competition is only a click away, an eBusiness can afford neither downtime nor a slowed network response time. According to Zona Research, an eBusiness must respond in less than seven seconds or risk losing 30% or more of prospective customers. Zona Research estimates that as much as $4.4 billion in e-commerce is lost each year because of failed or slow networks. Against this backdrop, Global Server Load Balancing (GSLB) has emerged as the new way of reducing the risk of losing your customers to the competition. GSLB improves web response time by transparently steering the customer to the nearest web server. GSLB improves web site availability by automatically redirecting the customer away from failed primary sites and to the nearest alternate site. Through these techniques, GSLB protects an eBusiness from natural calamities, catastrophes, or power outages and ensures that the customer gets to the nearest functional website.

Drawback of Standard DNS Operations

As illustrated in Figure 1, when a customer requests a URL such as http://www.foundrynetworks.com, their browser queries their local Domain Name Server (DNS) to resolve the host name www.foundrynetworks.com to an IP address. The local DNS performs queries

Figure 1: DNS Operation

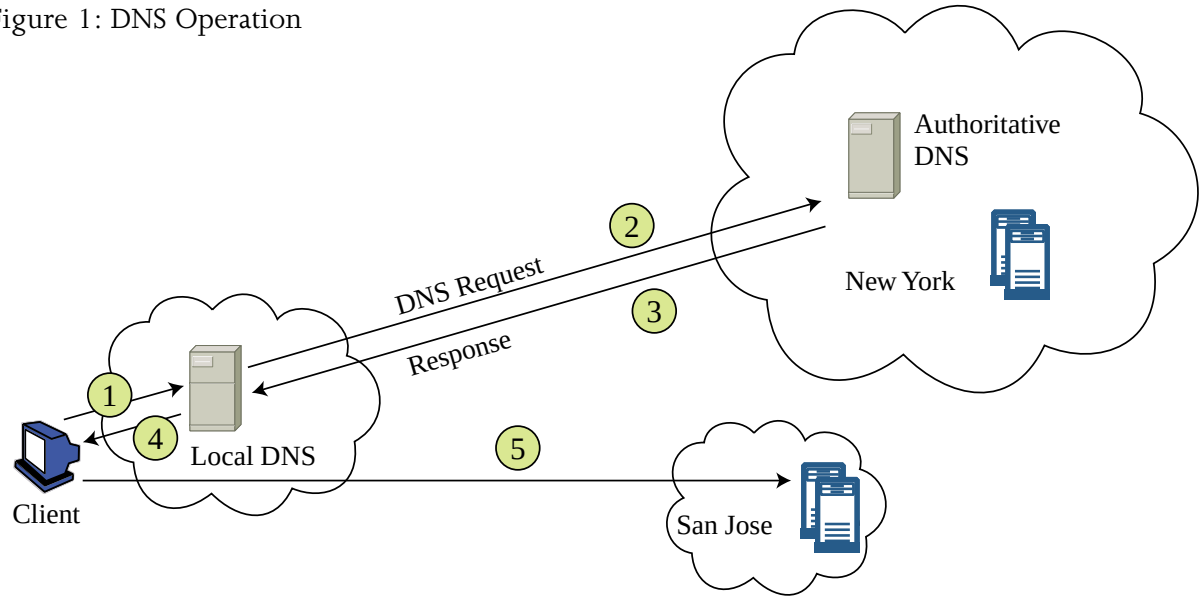

iteratively, ultimately reaching the authoritative DNS for the given domain foundrynetworks.com. The authoritative DNS replies with one or more IP addresses for the given domain name. Once the local DNS has the IP address, it responds back to the customer's browser, which in turns opens a TCP connection with the given IP address and downloads the web pages. The local DNS caches the authoritative DNS response and provides the same address for future requests to the domain until the time-to-live (TTL) parameter of the domain's IP address specified by the authoritative DNS expires.

DNS can provide a rudimentary form of load balancing if the authoritative DNS uses round robin across the IP addresses for a given domain.

There are four main problems with the above approach. First, if the server supporting the website is down, the authoritative DNS continues to return that server's IP address in response to local DNS queries until the authoritative DNS is manually updated. Second, even if the server is up and running, if the application that serves the website is down, customers receive an error code indicating that the requested website is temporarily unavailable, equally unacceptable to an eBusiness. Third, the authoritative DNS has no knowledge of proximity (where the customer inquiry originates). A customer accessing the site from Japan could be sent to a server in New York and vice versa, thereby burdening the customer with a slower response time. Fourth, the local DNS runs the risk of providing the customer with a bad IP address for the requested domain because the local DNS caches the authoritative DNS's original response until the TTL of the bad IP address expires – even if the authoritative DNS has updated information containing a new IP address.

ServerIron Global SLB

THE MOST INNOVATIVE ALTERNATIVE TO REPLACING DNS

ServerIron acts as a DNS proxy and becomes the front-end to the authoritative DNS. Optionally, ServerIron can also respond to DNS queries by itself. Because ServerIron combines server load balancing (SLB) and Global SLB into one convenient product, ServerIron load balances DNS traffic if there are multiple authoritative DNS servers. When the authoritative DNS replies to a DNS lookup request, ServerIron modifies the response to direct the customer to the best site available. ServerIron uses a sophisticated algorithm that consists of several policies and metrics in order to select the best site for a given customer. The GSLB ServerIron refers to the ServerIron doing global server load balancing and the site ServerIron refers to the ServerIron load balancing switch deployed at different sites for web contents.

ServerIron uses the following metrics to evaluate the site IP addresses in a DNS reply:

- Health conditions of each site, server, and application (Each site may consist of just one server or a group of servers load balanced by a ServerIron)
- The site ServerIron's session capacity threshold
- The round-trip time between the remote ServerIron and client networks
- The geographic location of the server and the client network
- The site ServerIron's available session capacity
- The site ServerIron's FlashBack speed (how quickly the GSLB ServerIron receives the health check results)
- The Least Response Selection (the site ServerIron that has been selected less often than others)

Health Checks and Flashback Speed

When acting as a DNS proxy to perform Global SLB for a given domain name, for example, www.foundrynet.com, the ServerIron makes a DNS look up request to the authoritative DNS to obtain the list of IP addresses that are serving the domain name. This means that you need to configure the list of IP addresses only once at the authoritative DNS. ServerIron makes periodic checks on the authoritative DNS to ensure the list of IP addresses is up to date. ServerIron performs layer 3, layer 4, and layer 7 (application) level health checks on each of the IP addresses. For this discussion, let's assume that each IP address in the DNS response represents a server farm load balanced by a site ServerIron. In addition to the health check, ServerIron also measures the FlashBack speed that represents layer 4 and application responsiveness. Because response times in the Internet vary by nature, ServerIron ignores minor differences in FlashBack speeds from different sites. The default tolerance level is 10%, and can be changed depending on the specific customer network requirements.

Geography based Site Selection

IP addresses are allocated in blocks to major geographical areas. By examining the IP address in a request, ServerIron can transparently steer customers to a site available in the same geographical area. In the example shown in Figure 2, ServerIron re-directs a client from Frankfurt to a site located in London, all else being equal. While this approach cannot determine the proximity of different sites to a customer if they all are located in the same continent, it ensures the request remains within the continental boundaries.

Site Proximity and Load Conditions

There are two major challenges in measuring proximity or round trip latency from the end user to each site. First, how can one measure response time accurately from an end user to each site? If a site ServerIron does a ping to the end user, the ping may be given

Figure 2: Geography based site selection

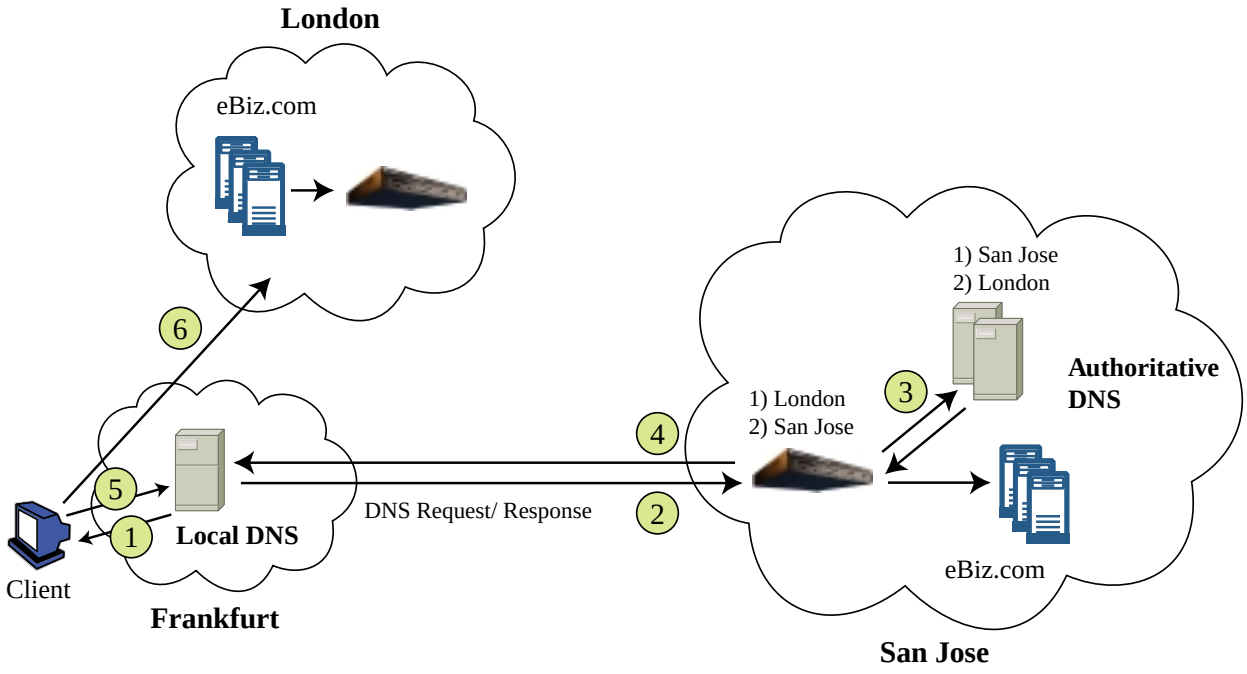

lower priority by routers or blocked by firewalls. If a site ServerIron attempts to establish a TCP connection with the end user, it may be blocked by firewalls or proxy servers. The second challenge is how to interpret the data. Even if the network response time was accurately measured, how can the global server load balancer use the response time data points? There can be hundreds of thousands or even millions of them. Further, no additional latency must be added at the DNS request in order to preserve the first impression in accessing a web site. ServerIron uses very innovative methods to solve both these challenges.

The GSLB ServerIron uses a proprietary protocol, the Foundry GSLB protocol, to communicate with the ServerIrons at the remote sites. The GSLB ServerIron uses the protocol to learn the following information from the site ServerIrons:

- The VIPs configured on the ServerIrons (Note that each ServerIron could be serving multiple hosts)
- Load Conditions
- Round trip time (RTT) to the customers accessing the site

The first time a customer accesses a web site, the GSLB ServerIron selects a site based on the available metrics. For example, the first time a customer from Denver attempts to access www.foundrynetworks.com, the GSLB ServerIron uses health checks, flashback speed, geography based site selection and load conditions to direct the customer to a site in New York. The customer's browser then opens a series of TCP connections to the site ServerIron in New York to download Web pages. The site ServerIron uses the natural traffic flow between the customer's browser and itself to measure the round-trip latency. Periodically, the site ServerIrons report the round-trip latency data points for its users to the GSLB ServerIron along with the site load conditions. The GSLB ServerIron aggregates all the data points into a proximity table indexed by network neighborhood. A network neighborhood is defined as a prefix of the customer's IP address. The default value for network neighborhood is /20, but can be changed to any value between /1 and /31. For each network neighborhood, the GSLB ServerIron can thus recognize the sites with the best round trip times. The ServerIron will ignore minor differences when comparing different round-trip time measurements. The default tolerance level is 10 percent and can be set to a different value based on specific customer requirements.

As more customers access the web site, the GSLB ServerIron builds a comprehensive proximity knowledge database that enables smarter site selection. In order to keep the proximity table useful and up-to-date, the GSLB ServerIron treats the proximity table like a cache, by freeing up the infrequently used entries to make room for useful data points. Because the response times in the Internet can vary over time, the GSLB ServerIron ignores the proximity metrics for 5 percent of the requests by default. This allows other metrics to direct clients to new sites to capture any changes in the response times. This is called exploration percentage, and can be changed depending on specific customer needs. The RTT measurements are independent of the domain name. When a site ServerIron measures RTT to a network neighborhood, this information is applicable to all domain names served by that ServerIron.

The GSLB ServerIron also uses the load conditions reported by the site ServerIrons in making the site selection. Because the GSLB ServerIron looks at available session capacity at each site, it allows each site to be built with different capacities. All else being equal, the GSLB ServerIron will prefer the sites with the lesser load. Each site can also be configured with are threshold value, and the GSLB ServerIron will avoid directing customers to those sites whose load has exceeded the configured threshold value.

Evolutionary Proximity Knowledge

ServerIron adds no latency to DNS lookups. It always applies the GSLB metrics based on the best available information at that moment. It then continues to collect various metrics from each site in the background and builds a knowledge base aggregated per customer network neighborhood.

Changes in Site Conditions After the DNS Lookup

Once the DNS lookup is complete, the customer's browser attempts to establish TCP connections with the given site's IP addresses. If there is a flash traffic surge that causes the given site to be suddenly overloaded or if all the servers or applications at the given site go down after the DNS lookup is complete, the site ServerIron can be configured to handle the requests in two ways. First, the site ServerIron can use HTTP redirect to send the customers to an alternate site. Second, the site ServerIron can load balance the requests across remote server farms.

Flexible Site Selection Policies

The GSLB ServerIron provides extremely flexible site selection mechanism. Depending on customer needs, a policy can be turned on or off, or the ServerIron can be configured to apply the site selection policies in a different order. For example, the GSLB ServerIron can

ServerIron GSLB provides the most comprehensive site selection

Site, server and application health checks	✓
FlashBack speed to measure site, server and application responsiveness	✓
Geography based site selection	✓
Site load conditions	✓
Configurable thresholds for site load conditions	✓
Site selection based on accurate proximity to customers	✓
Exploration to keep proximity tables up-to-date	✓
Handle server farm failures after DNS lookup	✓
Configurable tolerance values	✓
Ability to customize site selection policy	✓

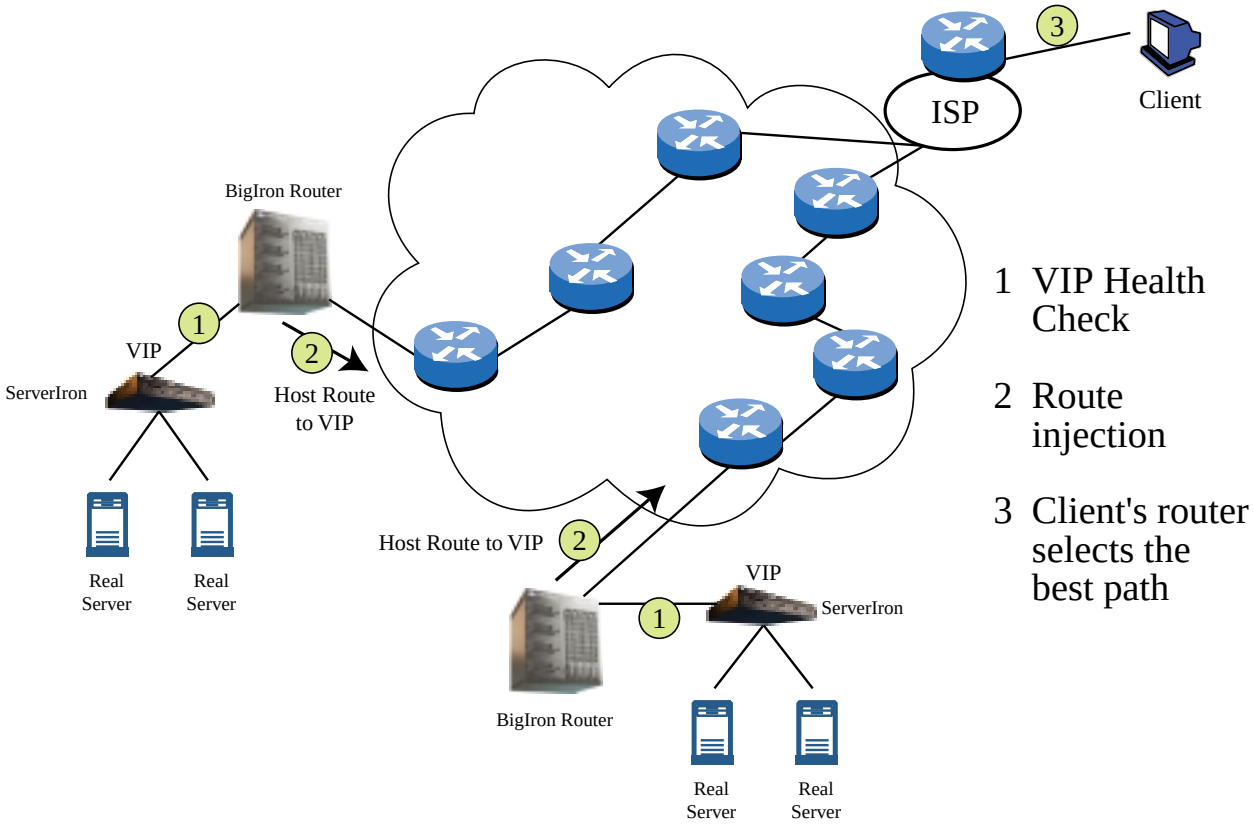

Figure 3: Global IP: GSLB using standard routing protocols

be configured to ignore geography-based information or to use site load conditions as the first metric in selecting a site.

Central Monitoring For All Site ServerIrons

Network administrators can review detailed statistics of site ServerIrons from the GSLB ServerIron. This approach minimizes the complexity in deploying and maintaining web sites and server farms at multiple locations.

Global IP: GSLB without using DNS

ServerIron Global IP provides GSLB without using DNS. Global IP provides one global IP address for web servers located throughout an ISP cloud serving a web site. It involves placing web servers for an eBusiness web site across the globe at different locations in the ISP's cloud connected to the ServerIron as shown in the Figure 3. Each geographic location can have multiple servers and can be load balanced with ServerIron. All ServerIrons service the same Virtual IP (VIP) address and are connected to a Foundry router such as BigIron or NetIron. The Foundry router checks the health of the virtual IP address on the ServerIron. If the health check is successful, the Foundry router injects a host route into the adjacent router in the cloud. If the health check fails, the Foundry router retracts the host route. When a customer attempts to connect to the virtual IP address of the web site, the ISP router near the client sees multiple paths to the same IP address. Based on the routing path cost, that router will select the best path, which leads to the nearest site.

ServerIron Overview

Foundry Networks' ServerIron family of Internet traffic management system switches provide high performance, Layer 4 through 7 switching, enabling network managers to control and manage today's exploding web transaction, web application and eCommerce traffic flows. Internet IronWare - Foundry Networks' unique software suite of Internet traffic management capabilities, powers ServerIronXL, ServerIronXL/G, and BigServerIron (a simple software upgrade to the BigIron chassis) to direct requests to the right server and application based on the information that resides beyond the traditional Layer 2 and 3 packet headers. ServerIron delivers industry leading performance for Internet traffic management functions including local and global server load balancing, firewall load balancing, transparent cache switching, application redirection, packet filtering and prioritization, and support for content-intelligent switching such as cookie-, URL-, and SSL Session ID-based redirection and load balancing.

Foundry's IronCore architecture, combined with custom packet processing ASICs, offers flexible deployment and support for extensive network topologies. ServerIron's shared memory architecture ensures exceptional concurrent connection capacity whether you use 2 ports or 24 ports. With an optional redundant power supply and a rack-optimized form factor, ServerIron provides the performance, port density, reliability, and flexibility required by every network manager and administrator.

2100 Gold Street
San Jose, CA 95002
Tel 408.586.1700
Fax 408.586.1900
www.foundrynet.com