Global Load Balancing

Open Enterprise Systems
Internet Technologies Group

Introduction

This document will discuss a technique, which Open Enterprise Systems (OES) has developed, to optimize geographically-distributed web servers. This technique fulfills two usage requirements:

1. **Speed.** Web visitors should be directed to the closest [1] available server for high-speed, low-latency service.
2. **Availability.** If a server experiences a problem (for example, software failure, hacker defacement, or upstream network disconnection), all further requests for service will be routed to the other servers.

The most effective path to meet these criteria is using DNS name resolution to dynamically return the IP address of the most appropriate server.

Smart DNS

Geographically distributing web services can be done using many methods, at various levels of the networking protocols. OES suggests using the DNS name resolution process to return the IP address of the correct server. Using DNS offers several advantages:

- **Simplicity.** All servers may be configured identically. The servers need not know about the existence of their counterparts elsewhere on the Internet.
- **Transparency.** DNS is much simpler than higher-level protocols, such as HTTP. The simplicity of a name resolution ensures 100% compatibility with all Internet standards.
- **Scalability.** Adding and removing servers is merely adding and removing IP addresses to the availability pool.
- **Upgradability.** Adding content, software, or even an entire server may be done on the Internet, but without adding the new IP address to the pool. This effects a testing environment which is identical to the production environment. When the service is ready, add the IP address to the availability pool.

Method

OES combines three major components to build the global load balancing service: mapping client IP addresses to desired server addresses, integrating this mapping into DNS software, and integrating service availability scanning into DNS software.

Location to Server Mapping

First, we build a database which maps Internet addresses to physical locations. In most cases, country of origin is sufficient, however, this system works with any resolution. Second, we build a mapping of geographical locations to the desired IP addresses which should provide service.

Integrating Mapping With Name Server

The mapping from the previous step must be integrated with the name server software[2] such that when an incoming request for name resolution arrives, the response is determined by the IP address of the requester.

Availability Scanning

Finally, if a server becomes unavailable for any reason, its IP address must not be used as a DNS response. Instead, the DNS software should return the next appropriate candidate. The DNS server itself determines which servers are available, and it maintains a pool of working servers. Availability means not only network connectivity, but also that the web content is as expected (i.e. nobody has hacked into the server and defaced the web site).

Example

Figure 1. Hypothetical Service Transaction

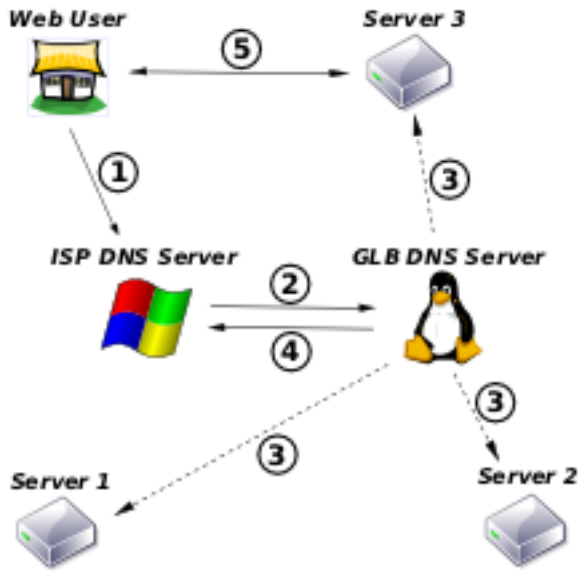

Hypothetical Service Transaction

1. The first time a user accesses the web site[3] , she contacts her ISP's DNS server.
2. The ISP's DNS server contacts our authoritative Global Load Balancing (GLB) DNS server and requests a DNS A-record (forward name) lookup.

3. The GLB server determins the best web server to handle the client's request by selecting the nearest server which is currently in the availability pool.
4. The DNS server returns the IP address of the closest web server.
5. The user begins a standard TCP transaction with the appropriate server.

Updating the Availability Pool: In practice, multiple DNS servers are used for redundancy; however, the structure of the DNS and TCP transactions is the same.

OES Service

Open Enterprise Systems provides this service to customers worldwide. We offer two approaches for our customers to use this system.

Managed Service

OES maintains multiple DNS servers (two by default, more at the customer's request) hosted at different Internet service providers with the requisite software already operational. After a brief collaboration period with our customers to determine server IP addresses and to calibrate our defacement detection components, we may immediately begin performing this service.

Delivery and Training

For customers who wish to have more control, OES will configure two or more new DNS servers explicitly designed to handle our customer's servers. Afterward, we follow the same process mentioned above, except on the customer's private servers. When the work is complete, we turn the entire system over to our customer's administrators, and we train their staff about the details of administering the system. Training is in-depth, including examining the source code modifications to the standard BIND DNS server.

Notes

1. For this discussion, the term *close* means close in the networking sense, i.e. the fewest hops along the lowest-latency route. For example, for servers based in Hong Kong and San Francisco, a user in eastern Russia is geographically nearer to Hong Kong. However, since Russia has more high-quality connections with the U.S., she should be directed to the U.S. for better performance.
2. This technique works with any DNS server, however OES uses BIND since it is already the primary carrier of domain name service and its source code is freely available.
3. In practice, users do name resolution through their ISP's DNS servers, and the ISPs cache the information locally. Thus if another user at the ISP makes the same request, her ISP will return the cached information. The ISP will only re-request the data once the DNS record has expired. However, this procedure is not fundamentally different from the example and only detracts from clarity.