DNS for Service Providers Part II - Availability

Frank Yue, 2013-15-01

It sure is no fun when I call a business only to get an automated message that says that this is not part of their normal business hours and I need to call again at a different time. It is even worse when the phone just rings forever and no one answers. Have you ever gotten a fast-busy signal when using your cell phone because all the circuits were busy or you did not have a good signal with your cellular provider? In today's fast-paced world, we expect everything to be available all the time. Imagine what it would be like if access to sites on the Internet were only available at certain times or if the DNS services were not available due to a hardware failure.

Communications Service Providers (CSP) are on the front line for support and access to Internet services. Not only do they connect customers' devices to the Internet, they provide the core services, such as DNS that are necessary for most applications to access needed resources.

In our previous post we talked about DNS scalability to be able to handle the potential load of traffic in case there was a surge in DNS requests and traffic hitting the DNS servers. Not only does the CSP need to provide the servers and the assisting technology such as DNS caching, they need to ensure that access to these resources is always possible and traffic can be distributed and directed to multiple locations to balance the load. Multiple DNS servers reside at each site and there are usually multiple locations for scalability and availability in case of a failure (power outage, natural disaster, human error, etc.).

DNS is Everywhere

There are several technologies that we will discuss and learn how they can and cannot provide continuous and reliable access to DNS infrastructure. The first is our tried and true server load balancing. CSPs will want to use a server load balancing technology to combine and balance traffic to multiple servers in a geographic location. Note that I restrict this to individual geographic locations since the traffic needs to be managed with the use of a virtual server which has to reside on a single device (or multiple in a high-availability scenario) in a restrictive physical site.

If an individual server or database fails, the server load balancer can detect the failure through the use of basic and advanced health checks and take servers out of the load balancing algorithm as well as put them in as necessary.

Some of you may be thinking that global server load balancing (GSLB) can alleviate that physical location problem. I would hope that many of you utilize GSLB technology today. But you need to remember that GSLB is essentially DNS response manipulation or load balancing via DNS. GSLB responds to DNS requests and provides the IP address of the best server and location for that particular DNS request. In our problem, we need to get the DNS request to the right DNS servers. Only then, can we apply GSLB technology to leverage a distributed resource architecture.

One technology growing in popularity to ensure that the DNS request is always able to reach a DNS server, even if a group of servers or an entire geographic location is unavailable, is IP Anycast. IP Anycast is the use of a single IP address for the DNS server and advertises that IP address from multiple locations and servers. It is necessary to use a routing protocol that supports IP Anycast such as OSPF or BGP. The routing protocol will direct the request to the best available location based on routing distance and metrics.

A combination of high performance server load balancing with the use of a geographic redundancy/availability protocol such as IP Anycast will ensure that a DNS request from a client will always be able to reach the DNS infrastructure. A proper DNS response can be delivered to the client and everyone will be happy with the 24/7 access to the Internet.

Note that the definition of availability can also include ensuring access even in the midst of a DNS DDoS attack. This will be discussed in the third segment on DNS for Communications Service Providers where I will talk about DNS security and ensuring the integrity of the infrastructure and data along with the management of DoS and DDoS attacks to deny access to DNS infrastructure and subsequently, the Internet.

F5 Networks, Inc. | 401 Elliot Avenue West, Seattle, WA 98119 | 888-882-4447 | f5.com

F5 Networks, Inc.
Corporate Headquarters
info@f5.com

F5 Networks
Asia-Pacific
apacinfo@f5.com

F5 Networks Ltd.
Europe/Middle-East/Africa
emeainfo@f5.com

F5 Networks
Japan K.K.
f5j-info@f5.com