Performance Measurements of DNS Root Service in China

NELNAT
April 24, 2015

When the Domain Name System (DNS) was originally designed, its global scope was not foreseen, and as a consequence only 13 root servers are deployed which provide the bootstrap foundation for the entire DNS system. The root principle to deploy 13 root servers at the beginning was that the limitation of DNS response message is 512 bytes.

As the Internet grew beyond its birthplace in America, academic community to span it increasingly put pressure on this limitation. Thus, anycast was presented as a solution since it would allow the system to grow beyond the static 13 instances, while avoiding a change to the existing protocol and root server placement. For a DNS root server, anycast provides a service whereby clients send requests to a single address and the network delivers that request to at least one, preferably the closest mirror site in that root server's anycast group.

One of the goals of anycast is to improve the resilience of the DNS infrastructure to Denial-of-Service (DoS) attacks. This problem cannot be solved simply by increasing hardware performance, because in most current deployments the servers can already withstand higher attack loads than the networks that surround them and during an attack it is network congestion rather than query load that renders the servers unresponsive. With anycast, DoS attacks are mitigated both by local nodes which act as local sinks for DoS attacks in their catchment areas, and by global nodes which spread the attack load over multiple servers and networks. The more service nodes deployed and the more widespread the deployment is, the less likely that an attack can cause widespread service disruption. Another goal of anycast is to improve performance. Deploying service nodes topologically close to clients will decrease query latency. Besides, anycast can increase the reliability of DNS service. Deploying nodes close to clients can increase reliability by decreasing the number of network elements that query message must traverse.

However, the fundamental reason of adopting anycast is that only 13 root servers were deployed but more and more DNS root serving instances are needed with the explosive deployment of Internet. Currently, 11 of the 13 letters are hosted at multiple countries, and the root zone is served at about 460 sites around the globe.

In order to illustrate what the situation of the mirror sites in China is, we deployed a monitoring system to measure the 7 deployed mirror sites in China from several aspects (it is anticipated that more and more countries or districts would share their measurements about the local[1] mirror sites in the near future[2]) and these mirror sites are illustrated in Table 1.

[1] The term "local " here does not illustrate the service scope of the anycast node, but means that there is a DNS anycast node in this area and this node may be globally or locally configured.

Table 1. DNS mirror sites deployed in China

Mirror site	Operator	Location	Deployment year
F	China Telecom	Beijing	2003
I	CNNIC	Beijing	2005
J	China Unicom	Beijing	2006
L	CNNIC	Beijing	2012
L	ZDNS	Beijing	2014
L	BII	Beijing	2014
L	VNET	Beijing	2014

This monitoring system consists of 61 monitors in 32 provinces, covering 6 main ISPs in China (mainland). Figure 1 shows the averaged resolution latency of the 13 root servers from two main ISPs in China (In consider of the business secret protection, the company title of the two ISPs are not given here because it is inappropriate to show them together) and the time duration is from 10th April 2015 to 15th April 2015. Generally, the resolution performance of the root servers with local mirror site has appeared to be better while the diversity is obvious. However, even with the local mirror site, the performance also may be affected by many factors, for example, the F root service in ISP2 bares higher latency even compared with some servers without mirror sites in China (maybe affected by the network environment of ISP2). The results suggest the instability of the root service, considering the TLD resolution was affected significantly when the local mirror site or its corresponded root server fails.

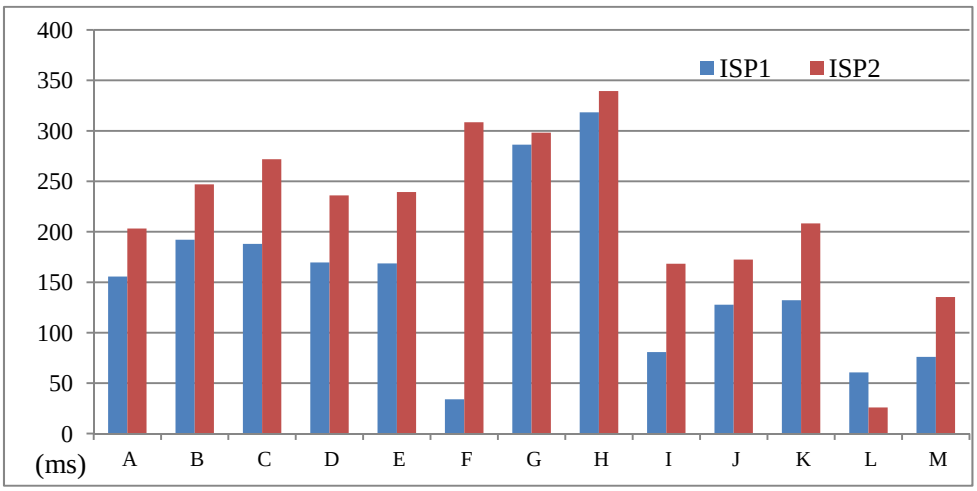

Figure 1. Averaged root server resolution latency

It shows that the root servers with local mirror sites in China (F/I/J/L) have lower latency compared with others. In the following, we take two servers to illustrate their performance in different locations. Figure 2 shows the averaged latency of the F root server from different provinces.

[2] Different monitoring system may produce different results due to different number of monitors, their locations and configurations.

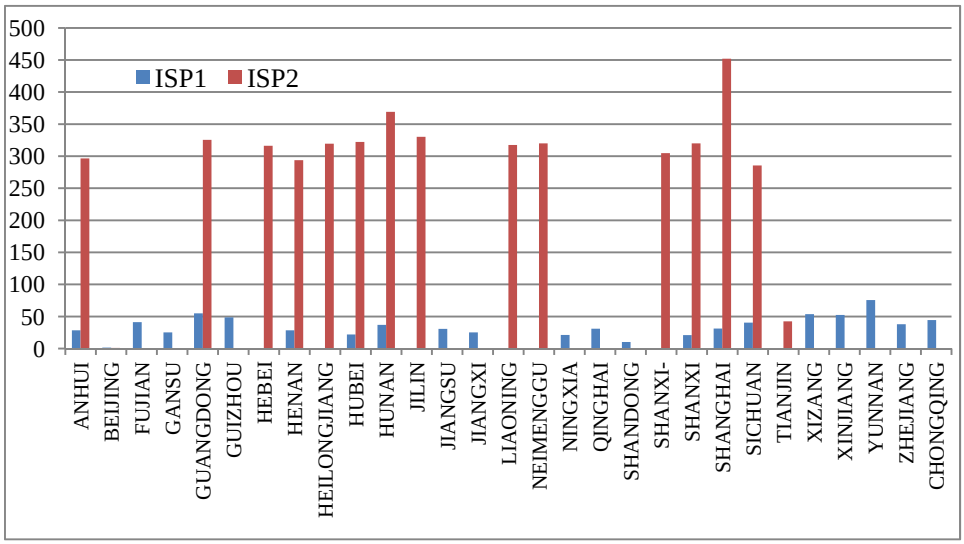

Figure 2. F root server resolution latency

The results of two ISPs are mixed in Figure 2 (and the following figures) because the monitors do not cover both ISPs in all the provinces. For example, we only have a monitor in ISP2's network in HEBEI province. The results indicate that the latency in different locations has huge difference; most of them in ISP1 are under 50ms while the latency in ISP2 exceeds even 450ms in SHANGHAI. For ISP2, there may be some network outages on the routing table configuration and link reachability, while for ISP1, the basic reason is that most of the resolutions to the F root server hit the F mirror site in China as shown in Figure 3.

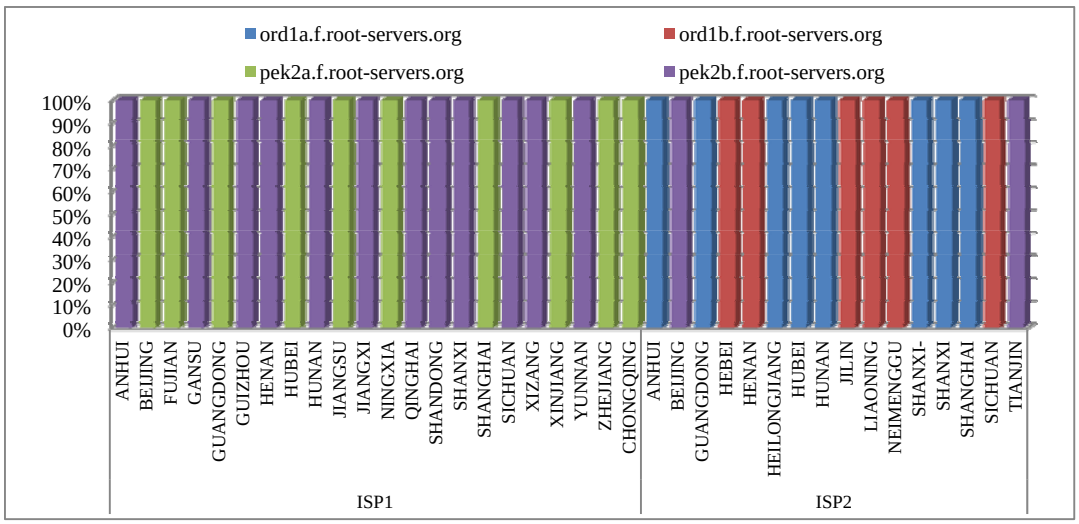

Figure 3. Hit ratio of F mirror site in China

In Figure 3, "pek2a" and "pek2b" denote two IP addresses (two servers) used by the F mirror site in China. "ord1a" and "ord1b" denote two IP addresses (two servers) used by the F mirror site in America (Chicago). As shown, the 10 measurements to the F root server from ISP1 all hit the mirror site in China, while the measurements to the F root server from ISP2 all go to the mirror site in America and that causes the prolonged latency to F root server from ISP2.

Then what about other mirror sites? Is ISP2 always so badly serving? Figure 4 shows the averaged latency of the L root server from different provinces.

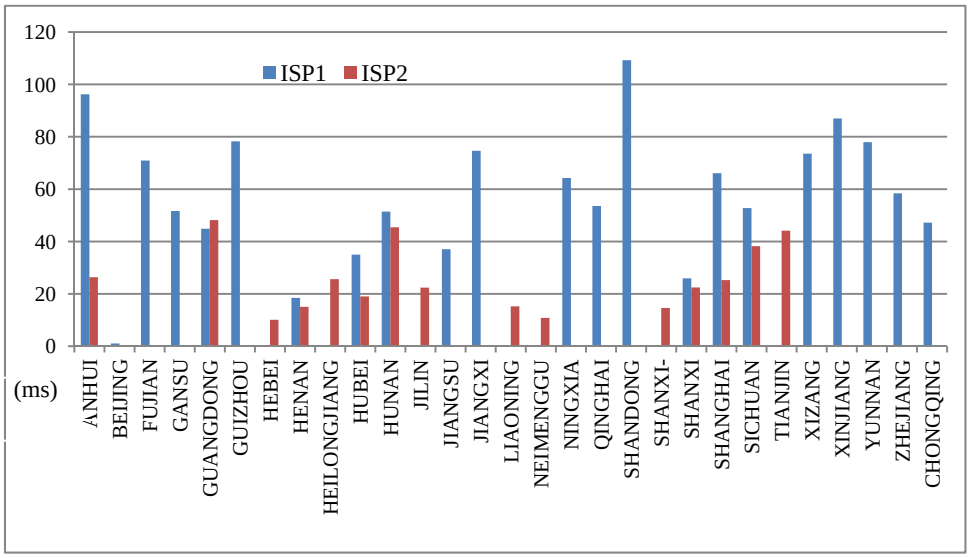

Figure 4. L root server resolution latency

Compared with the performance of F mirror site in ISP1, the averaged latency of L root server is significantly higher. While for ISP2, the situation turns better. Accordingly, the hit ratio of the L mirror site in China is shown in Figure 5.

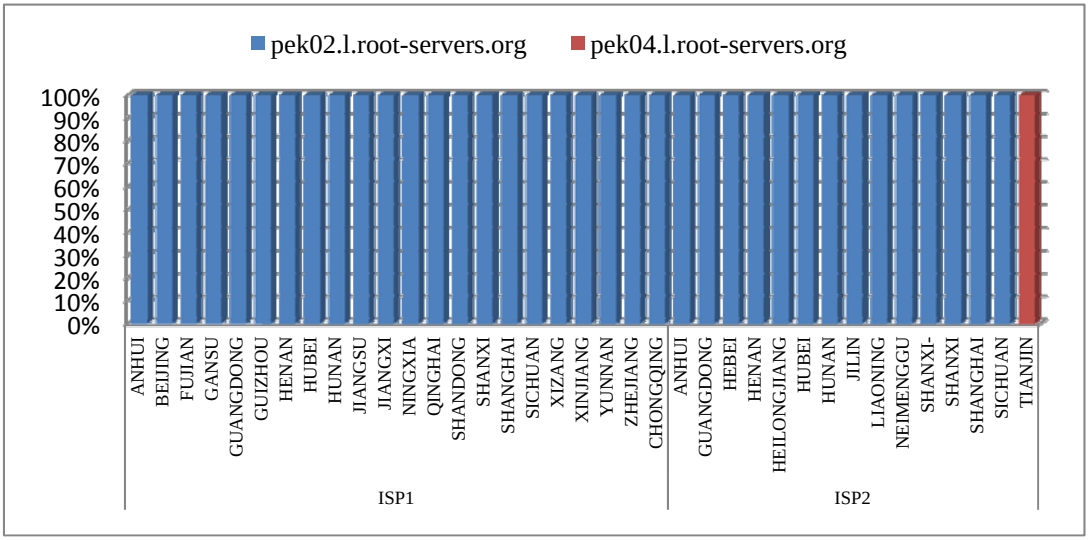

Figure 5. Hit ratio of L mirror site in China

In Figure 5, “pek02” and “pek04” are two IP addresses of L mirror site in China. As shown, the 10 measurements to the F root server from two ISPs all hit the mirror sites in China. Even for ISP2, its requests are served by the local mirror sites and then the performance is improved significantly.

Without doubt, these results are influenced by multiple factors, such as:

- ✧ Numbers of recursive servers deployed in each province are different;
- ✧ ISP may have diverse routing policy for the DNS resolution;
- ✧ Different locations and ISPs bare different bandwidth and traffic load;
- ✧ The distance from each province or ISP to the mirror sit is different.

Based on these factors, during our measurements, we also found that the requests may be automatically redirected to the other mirror sites even the local site is deployed.

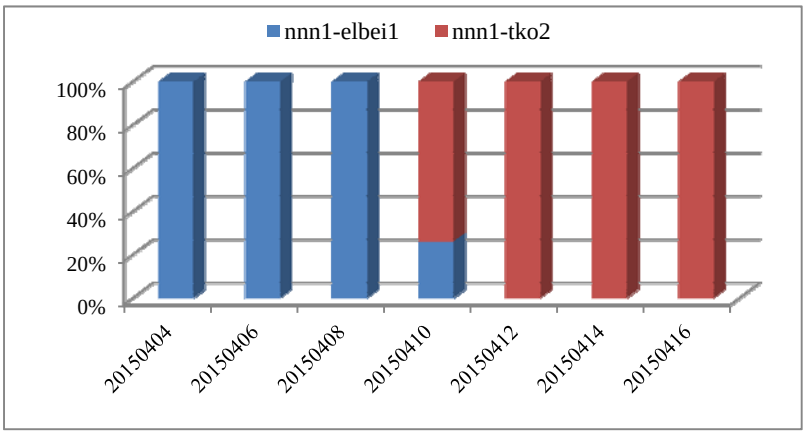

Figure 6. Redirection between the mirror sites

An example is shown in Figure 6, before 10th April, requests from minitors in SHANGHAI in ISP1 to the J root server were all served by the mirror site in China (e1bei1), while afterwards, the requests from the same minitors were redirected to the mirror site in Japan (tko2).

The performance of the DNS root service is very complex and affected by many factors as exampled above. In the future, we will continue this work from the two following motivations:

1) To discover the network problems causing the performance degradation of the DNS root service in China.
2) To figure out the optimized deployment policy of DNS root service in China from the long-term measurements.

Authors of this report:
Zhiwei Yan
Email: yanzhiwei@cnnic.cn
Lanlan Pan
Email: panlanlan@cnnic.cn