How to scale the DNS root system? draft-lee-dnsop-scalingroot-00.txt

Xiaodong Lee

Paul Vixie

Zhiwei Yan

@ IETF 90, Toranto

Background

- Initial 13 root servers should be scaled from many aspects:
  - Academical and industrial researches about their uneven distribution and possible solutions to scale
  - More and more CDOs (Country/District/Organization are eager to hold a root server to serve their local regions.
  - Both ICANN and IETF start this work!!!

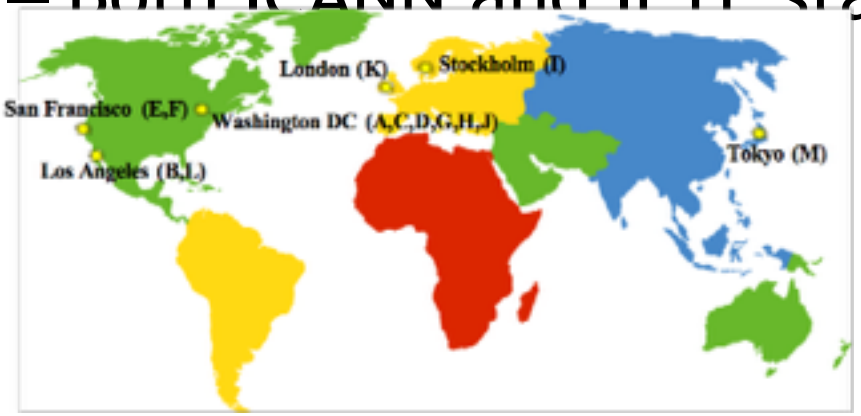

ICANN

IETF

Motivation

- 512-byte limitation was broken due to IPv6.
- More root servers (N) can be deployed under IPv6:

DNS Priming Response Message (IPv4 and IPv6)	#Bytes
Required Header	12
Query	5
Answer: first answer is 31 bytes, the sequent answers are reduced by 16 bytes due to compression	31+15(N-1)
Additional Records : each A record is 16 bytes	16N
Additional Records : each AAAA record is 28 bytes	28N

$$12 + 5 + 31 + 15 \times (N - 1) + 16 \times N + 28 \times N = 1232$$

  - N=20. Who should hold the 7 more root servers?
  - It is a new headache. We need a real solution!!!

Possible Architectures

- With DNSSEC, DNS data is trustful no matter where I got it!---*it is a exciting basis.*

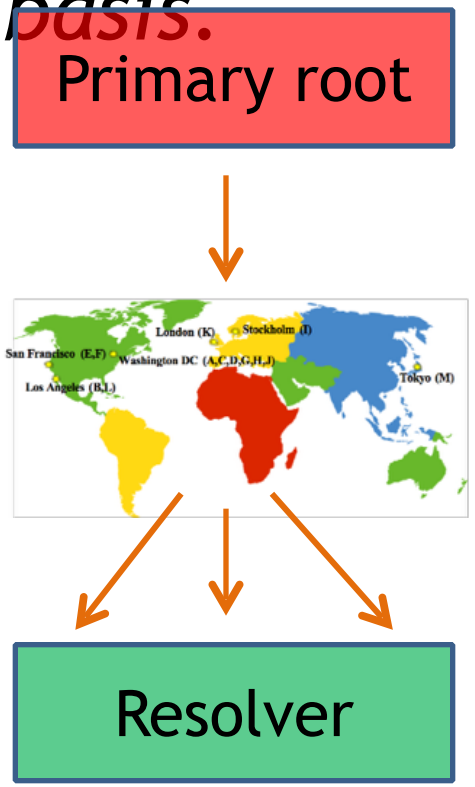

- **Where should we start?**
  - Authoritative-level
  - draft-lee-dnsop-scalingroot
  - Recursive-level
  - draft-wkumari-dnsop-dist-root

Our architecture

- We can copy from both primary root server or one(multiple) slave root server(s).
- The common idea:
  - Anycast address is globally used by the CDO, who has the interest and ability to hold a root server.
  - Root zone file is synchronized under the DNSSEC protection.
- Difference: allocate new anycast addresses or use the current anycast addresses.

Solution 1

- Allocate new anycast addresses (for root server X):

  . IN NS anycast-X1.iana-servers.net.

  . IN NS anycast-X2.iana-servers.net.

  $ORIGIN iana-servers.net.

  anycast-x1 IN AAAA 2001:?:1::1

  anycast-x1 IN A ?.?.1.1

  anycast-x2 IN AAAA 2001:?:2::2

  anycast-x2 IN A ?.?.2.2

- These NS records will denote name servers whose addresses are not assigned to any current RNSO but are instead held in trust by IANA for use by any or all interested CDOs.

Solution 2

- Use the current anycast addresses:

  One or multiple root server addresses (within the 13 root servers) should be used as the anycast addresses of CDO DNS root servers.

- The CDO root servers in this case may be global node or local node.
- The CDO can manage and control this node (such as copy mirrors) according to their own needs.

Data synchronization

- Different synchronization may be used:

Distribution master	Solution1	Solution2			
IANA(-like org.)	√	*Global*	√	*Local*	
RNSO			√		√

- *A more scalable scheme is needed if the CDO root servers increase rapidly.*

DNS request to root server

- The DNS request to root server may be served anywhere. For example:
  - the request message originated from the edge user is served by the root server deployed in the local network within the same ISP.
  - the request is transmitted to the nearby ISP due to the routing protocol and responded by the root server in the nearby ISP.
  - the request may travel all the way to the Internet backbone without having been answered closely, in which case it may be answered by an RNSO root server for example.
  - ……

Difference with Warren's solution

- Resolver can of course caches the root zone file and serve its requesters efficiently.
  - This is a god approach especially for the large resolver (such as Google's open resolver).
  - However:
    - 1) the root zone file should be synchronized in an efficient and scalable manner.
    - 2) deployment difficulties (Recursive resolver= Recursive resolver + root server).
    - 2) it can not fulfill the requirements of some CDOs.
  - Anyway, they are different solutions and suitable for different scenarios and requirements.
  - ……

My suggestions

- Step 1
  - discuss their feasibility from technological aspects.
- Step 2
  - figure out respective scenarios and requirements suitable for each of them.

Thank all of you!