DNSOP
Internet Draft
Intended status: Informational
Expires: May 30, 2015

Di Ma
Feng Han
ZDNS
Linjian Song
BII

November 26, 2014

Considerations on the evolution of DNS root zone operation and management
draft-madi-dnsop-drzom-00

Abstract

Given responsibilities and importance it bears, DNS root system architecture remains largely unchanged. Over the years, the topics together with some proposals of scaling the root have been discussed in various communities, trying to offer solutions or insights to a specific issue of DNS root operation. This document gathers and describes issues relating to the root zone operation and management and comb corresponding technical requirements for evolution directions of root zone operation and management in new groundwork.

1. Introduction

DNS service links naming space and addressing system, a connecting thread through the Internet infrastructure. DNS was designed as a hierarchical system with a few TLDs and slow update frequency. Rarely has the domain name system faced so many significant changes in its root zone operation and management, with DNSSEC-signed zone file, IPv6 adoption and new gTLD. Besides, DNS is expected to deliver more responsibilities by supporting DANE (DNS-based Authentication of Named Entities) that allow Internet applications to establish cryptographically secured communications by using information made available in DNS, which by the way, has been launching discussions in IETF DANE WG. Given responsibilities and importance it bears, DNS root system architecture remains largely unchanged. Issues of root security and stability have been stated in many technical reports [RFC2870] [SAC042] [SAC046]. Over the years, the topics together with some proposals of scaling the root have been discussed in various communities, trying to offer solutions or insights to a specific issue of DNS root operation.

"Scaling the root" was one of remarkable issues when it comes to evolve the current root system. It could be to extend the number of root name servers, or it could be to make the root operation mechanism scale to bigger zone file and higher update frequency, or it could still be to establish a more open root zone publication system. DNS root operation is critical service to the Internet and any changes to it ought to be in a prudent way, assuring security and stability. Among others, zone file management, the focal thing in root operation, is a systematic task, involving kinds of multi-stakeholders in Internet community. We cannot get the pleasing outcome for root scaling without reviewing the DNS root operation and management on the whole.

This document, inspired by [icann-ITI] and [RSST], is intended to gather and describe issues relating to the root zone operation and management and comb corresponding technical requirements for root zone evolution directions, before the IETF community gives solutions to the very issues.

2. Modeling DNS root system operation

2.1. DNS root system components

As described in many technical reports and also a common view, there are three sub-systems involving DNS root zone operation and management:

The Resolving System publically interfaces with DNS resolvers, by responding to root zone query requests via DNS protocol data units.

The Publication System distributes the root zone file to the root servers (including anycast clones) and arranges for the information in the root zone file to be available for the construction of responses to queries.

The Provisioning System processes change requests, maintains the authoritative database of root zone information, and periodically creates a root zone file that contains the authoritative root zone information.

2.2. Roles in DNS root system operation

In accordance with this model, DNS root system roles could be therefore identified:

Root Zone Coordinator (RZC): RZC works as an administration body of root zone operation. RZC, as a workflow interface with TLD registries, receives and processes root zone change requests and executes administration to TLD registries. With DNSSEC deployment, RZC additionally delivers the responsibility of root zone Key Signing Key (KSK) management. Accordingly, RZC is recognized as a root zone operation representative to the public. The role of RZC is part of IANA functions, acted by ICANN currently.

Root Zone Host (RZH): The role of RZH is to maintain reliable, secure, and accurate operation of the DNS name servers that publish root zone data they fetch from RZC (IANA). As a distribution channel for root zone data, RZH is responds to root zone information queries via DNS messages. As well known, there are 13 sets of root name server IPv4/IPv6 addresses reflecting hundreds of individual machines representing RZH, denoted by 13 letters from A to M, operated by 12 independent organizations worldwide.

Root Zone Maintainer (RZM): RZM is the backend operator of root zone. Viewed as a technical role, RZM takes the responsibility of generating (updating) root zone file together its digital signature, and distributes DNSSEC-signed zone file to Root Zone Host (RZH). RZM.

At present, VeriSign serves as RZM.

Root Zone Administrator (RZA): RZA is privileged to perform the review and approval function that authorizes each individual change to the root zone. The RZA imparts nothing to the root zone but merely explicitly authorizes these changes by verifying that RZC has followed established policies and procedures in processing the requests. At present, the role of RZA is played by National Telecommunications and Information Agency (NTIA), a governing body of the United States.

3. Evolution directions

This section is to list the evolution directions with the purpose to make DNS root system remain an unshakable corner stone of DNS operation with security, stability as well as high efficiency, while provide a better and up-to-date service. Responding to root queries with security and stability is the fundamental among others for root system. In this document, the evolution of DNS root system is identified hereby as a set of key issues with respect to resolving system, publication system and provisioning system.

3.1. DNS protocol data unit and Resolving System

One of the fundamental problems of DNS protocol is the UDP size limitation, which is caused by the IPv4 MTU setting in the early days of Internet. Efforts made by EDNS introduce a negotiation mechanism to support bigger DNS protocol data unit. However, owing to either capability or filtering policy of intermediate networking devices, EDNS fails to get the pleasing expectation that it is designed for. Granted, EDNS is improving over time with updated versions. IPv6 is yet another thing that DNS protocol data unit can benefit from. IPv6-capable mandates an MTU of 1280 bytes, which means, even without EDNS, the size of DNS protocol data unit could be still improved by the virtue of IPv6 fundamentally. Where and when acceptable, TCP is still another option for carrying DNS messages [RFC5966].

With a bigger DNS protocol data unit, more information can be included in DNS message, especially offering a chance to enhance the Resolving System. By far, there are significant issues regarding the Resolving System owing to the limitation of the size of DNS protocol data unit.

When a recursive name resolver is bootstrapped, it uses a hints file or other statically pre-configured initial glue to find a root server, and then it asks that root server for the current list of root servers, with the expected answer the full list of RZH and their addresses, the process of which is called "priming exchange". For one

thing, with the inclusion of IPv6 addresses for RZHs, the response is even longer. In the case of 13 RZHs, there is no placing all the IP addresses of all dual-stake RZHs in a priming response with UDP limitation of 512 bytes. For another thing, it has been decades since there were 13 RZHs, which was a calculated one regarding to specific constraints. The number of RZH is by no means a constant for DNS system but a variable parameter when and where possible. And still for another thing, when DNSSEC signatures are added to the root zone, the response message to the priming query or DNSKEY query will exceed 512 bytes accordingly.

Together with the tendency that DNS is delivering more responsibilities, bigger protocol data units will therefore facilitate advancements designed to scale the root system in the days to come.

3.2. Root zone synchronization and Distribution System

DNS root sub-systems are not organized closely to one another but root zone keeps changing. Data synchronization is therefore fundamental to root zone distribution from RZM to RZH via Distribution System. Considering DNS tends to be more dynamic, synchronization should be taken into account when a new Distribution System is introduced.

In alignment with current DNS operation, DNS root name servers ({A-M}.root-servers.net), serving as RZH, receive queries from clients using the DNS protocol and provide appropriate responses. As one important channel to publish root zone information, the root name servers have been subjected to attacks over the decades, mostly of the Distributed Denial Of Service (DDOS) variety. Besides, reachability of the root name server system is required even for purely local communication, since otherwise local clients have no way to discover local services. As a point of view from [icann-ITI], in a world sized distributed system like the Internet, critical services ought to be extremely well distributed.

Historically and from the perspective of protocol, the channel for distributing root zone information guarantee the authenticity. This situation has been changed by DNSSEC that was designed to provide a security extension to DNS. One of the ways that we can benefit from DNSSEC is DNSSEC-signed root zone file itself guarantee authenticity no matter where it is fetched as long as contents of this zone file could go through the validation based on DNSSEC information with a globally-trusted single key maintained by IANA. Therefore, the deployment of DNSSEC and the advantage it brings about could evolve the root zone Distribution System scaling to a more open fashion. Accordingly, an open publication system of DNS root zone is

desirable. The role of RZH is expected to be played in a different way. By the virtue of DNSSEC-signed root zone, theoretically, any entity with good credit and network could be authorized to host root zone as long as the hosted zone file is approved and signed by IANA. Note that the selection process of the RZH candidate is out of the scope of this document.

Some efforts have been made. For instance, [I-D.dnsop-dist-root] proposes an enhancement that recursive DNS resolvers, serving as RZH, get copies of the root zone, validate it using DNSSEC, populate their caches with the information, and also give negative responses from the validated zone. While [I-D.dnsop-scalingroot] positions that IANA produce several additional forms of the DNS root zone by creating yet another "golden address" pair. It asks the IANA to authorize an un-owned pair of addresses that anyone can hang root service on.

Although promising, there are problems gripping the designs for open publication system of DNS root zone. If DNS root zone publication system operates in an open mode, far more RZHs are expected to take part in. Given an increasing number of RZH, either recursive resolvers or "universal anycast" endpoints or anything else, it should be taken into account how to scale the zone file synchronization mechanism to the much bigger groups of RZHs, ensuring consistence and negligible delay. Optimized root zone file distribution mechanism would be desirable. For example, TLD KSK rollover should be accomplished within a specific time, requiring the DS records update should synchronized beforehand timely.

Besides, open publication brings about a more diverse group of RZHs, including ones that are deployed in a poorly-connected Internet locations. And an increasing size of root zone file can easily be served from sites with high-bandwidth connections and ready access to servers and other computing infrastructure. It cannot easily be served from sites with poor connectivity or infrastructure. A synchronization mechanism of this open publication that can scale to networking diversity is also indispensable.

3.3. Root zone update and Provisioning System

3.3.1. Update frequency

The size of root zone file is attached with significance. "Size" here is not merely referred to as how many bytes the zone file has, but also with the number of delegations considered. A conspicuous observation is root zone has been growing up. As analyzed in [icann-impact-newgTLD], the number and size of records in the root zone has successfully grown over time, in part to accommodate new developments like the introduction of IDN ccTLDs, the first two rounds of new

gTLDs, the introduction of IPv6 glue records, and the deployment of DNSSEC in the Root Zone. There is also a natural tendency for the number of name servers per delegation to increase as TLD name server infrastructure matures over time.

As indicated in [RSST], due to introduction of DNSSEC and TLD update, the transferring frequency of root zone will, as estimated and desirable, remarkably increased, bring about the undesirable latency, which will place a significant impact on root operation. The very observation and analysis is yet another significance that should be included into the context of root zone management evolution.

DNS root zone is getting bigger in an unprecedented rate and to an unpredictable size. Ever since the new gTLD program was launched by ICANN, more than 400 new gTLDs have been delegated so far. As a natural tendency and estimated, more and more TLD operators will get involved in root system, thus the rate of change or update would increase proportionately. These changes includes:

1. Delegation and re-delegation
   These are additions of new top-level domains or the transfer of operation of a top-level domain from an existing operator to a new operator. There might also be the removal of a top-level domain.

2. Changes in contact information
   There are usually three official points of contact for a top-level domain, the formal head of the operator, the administrative contact and the technical contact. Each of these can change from time to time. The Root Zone Registration Data ("WHOIS") system is therefore affected.

3. Changes in the set of name servers
   Each top-level domain is served by two or more name servers. Top-level domain operators occasionally add to, change, or remove name servers from their set.

4. Changes in the addresses of name servers
   Name servers are occasionally renumbered. Also, when a new name server is added to the set serving a TLD, its address must also be added.

5. TLD KSK rolling
   TLD operators request to update their DS records in root zone.

   For each change, current root server system has several parties involved which including RZA, RZS, RZM and RZHs. And the whole process consists of several manual steps, which dramatically

decrease the efficiency and predictability of the update process. To make the whole process automated and use other zone update policy (for example, dynamic update), measures should be taken to fulfill the impending improvement to the Provisioning System.

3.3.2. Shared zone control

Over the years, concerns have been raised towards current root operation is based in the Untied States and subject to the jurisdiction of the United States. Due to the multi-stakeholder mode advocated and coordinated by ICANN, shared control on root zone is desirable by many and technical specification is needed therefore. One straightforward thought is that root zone data should have multiple signatures. Some theories have been advanced together with multiple signing protocols.

Technically, efforts have been made on coordinating DNSSEC signing information and multiple signing protocols. But how to make these technologies fit into the context of shared control on DNS root is still an open question, which should follow current root operation practices with new risks evaluated. As indicated in [icann-ITI], the right model is one in which all of the parties sharing control have a set of capabilities:

1. A system for initiating a shared zone consisting of the zone itself, rules, and individual journals for each of the participants to post their requests and actions.

2. Each type of request is visible to all of the other participants who can approve, disapprove, or timeout.

3. Rules define what happens to a request
   * One type of a rule is a vote, which defines the conditions for a request to succeed. This might include a delay for all parties to have time to consider the request. For ccTLDs the WSIS rules would dictate 1 of N, so each Country Code Top Level Domain (ccTLD) could unilaterally change its own data. Other domains might use a simple majority

   * Specified delays could be important so that others might be able to point out operational issues and let the requesters reconsider

   * Different conditions might apply for different operations, such as creating a new vs. editing, etc.

4. Security Considerations

TBD

5. IANA Considerations

TBD

6. Acknowledgements

The authors would like to thank Bill Manning and David Conrad for reviewing this document.

7. Informative References

[RFC2870] R. Bush, D. Karrenberg, M. Kosters, R. Plzak., "Root Name Server Operational Requirements", RFC 2870, June 2000.

[SAC042] SSAC Comment on the Root Scaling Study Team Report and the TNO Report, https://www.icann.org/en/system/files/files/sac-042-en.pdf

[SAC046] Report of the Security and Stability Advisory Committee on Root Scaling., https://www.icann.org/en/system/files/files/sac-046-en.pdf

[RFC5966] R. Bellis., "DNS Transport over TCP - Implementation Requirements", RFC 5966, August 2010.

[icann-ITI] Identifier Technology Innovation Panel, https://www.icann.org/resources/pages/identifier-technology-2013-10-11-en

[I-D.dnsop-scalingroot] Xiaodong Lee, Paul Vixie and Zhiwei Yan., "How to scale the DNS root system", draft-lee-dnsop-scalingroot-00(work-in-progress), July 2014.

[I-D.dnsop-dist-root] W. Kumari, P. Hoffman., "Securely Distributing the DNS Root", draft-wkumari-dnsop-dist-root-01(work-in-progress), July 2014.

[icann-impact-newgTLD] Joe Abley and Kim Davies., "Impact on Root Server Operations and Provisioning Due to New gTLDS", June 2012.

[RSST] Jaap Akkerhuis, Lyman Chapin, Patrik Faltstrom, Glenn Kowack, Lars-Johan Liman and Bill Manning., "Report on the Impact on the DNS Root System of Increasing the Size and Volatility of the Root Zone", September 2009

Authors' Addresses

Di Ma
ZDNS Ltd.
4, South 4th Street, Zhongguancun
Haidian, Beijing 100190
China

Email: madi@zdns.cn

Feng Han
ZDNS Ltd.
4, South 4th Street, Zhongguancun
Haidian, Beijing 100190
China

Email: hanfeng@zdns.cn

Linjian Song
Beijing Internet Institute
2508 Room, 25th Floor, Tower A, Time Fortune
Beijing 100028
China

Email: songlinjian@gmail.com