DNS Root Service Operations

ICANN Office of the Chief Technology Officer

Paul Hoffman

OCTO-021v2

23 March 2021

ICANN

TABLE OF CONTENTS

This document is part of ICANN’s Office of the Chief Technical Officer (OCTO) document series. Please see the OCTO publication page for a list of documents in the series. If you have questions or suggestions on any of these documents, please send them to octo@icann.org.

This document supports ICANN’s strategic goal to strengthen DNS root server operations governance in coordination with the DNS root server operators. It is part of ICANN’s strategic objective to strengthen the security of the domain name system and the DNS root server system.

The author would like to thank the many people who contributed to this document. The first version, OCTO-021v1, had substantial input from Matt Larson, David Conrad, Terry Manderson, and Mauricio Vergara. The second version had useful suggestions from Steve Crocker, Kim Davies, Wes Hardaker, Ted Hardie, Ken Renard, and Duane Wessels.

1 Introduction

The root zone of the global DNS is the starting point for all DNS data. Access to the DNS root zone is vital to Internet users and the primary way it is made available to resolvers is via the root server operators (RSOs). The RSOs are independent from each other, but they agree to a common set of eleven principles, laid out in Section 3 of RSSAC037, “A Proposed Governance Model for the DNS Root Server System,” for how the root server system (RSS) should be managed. Here is a summary of four of the primary principles that speak to the heart of operations of the RSS:

- IANA is the source of DNS root data, and RSOs are committed to serving the IANA global root DNS namespace.
- The Internet Engineering Task Force (IETF) defines technical operation of the DNS.
- RSOs must operate with integrity and an ethos demonstrating a commitment to the common good of the Internet.
- RSOs must be neutral and impartial in their service of the root zone.

The impact of these principles on root service operations cannot be understated as it drives the ethics of operations for the RSS’s important function. The eleven principles in RSSAC037 form the core of RSO’s commitments, many of which go well beyond technical operations.

Many ICANN participants who understand how some parts of the overall DNS operate may not be fully aware of how the RSS is managed. This document is aimed at readers who understand a reasonable amount about the DNS but who have not yet studied the operations of the RSS. Note that the document covers the operations of the whole RSS, not the operations of an individual RSO. It shows that there are numerous operational choices in most if not all of the parts of the service to make the overall service highly available and reliable for handling the current level of approximately 100 billion DNS queries a day. It also shows that there is no single best choice for many of the options.

For DNS name service, the field of operations encompasses many areas including:

- Being able to respond to a high volume of DNS queries, including during occasional but massive denial-of-service (DoS) attacks
- Having multiple systems in order to respond to DNS queries more quickly or to have greater resiliency
- Responding to queries sent from any valid IP address
- Handling normal and emergency Internet routing shifts
- Correctly implementing IETF standards and being responsive to DNS protocol evolution

Additional operational and security requirements for the RSOs can be found in RFC 7720, “DNS Root Name Service Protocol and Deployment Requirements."

Operations for the DNS root service are steeped in history. This document explicitly does not delve into the history of the choices that affect current operations. The Root Server System Advisory Committee (RSSAC) published an in-depth history of the root server system, RSSAC023v2, “History of the Root Server System.” It includes descriptions of each RSO’s service, from that RSO’s perspective. Section 5 of this document talks about how operations of the root server system are expected to change in the future.

RSSAC has created its own vocabulary, much of which is relevant when discussing operations. RSSAC026v2, “RSSAC Lexicon” is a concise source for understanding these terms; of particular interest in that document for understanding RSS operations is the definition of “instance.” Other DNS-related terms come from RFC 8499, “DNS Terminology.”

2 Root Zone Distribution to the RSOs

Root service operations, as described in this document, starts at the point that the root zone maintainer (RZM) creates a new version of the root zone based on direction from IANA. The RZM function is contracted by ICANN; the contractor is currently Verisign. The RZM publishes the root zone on a system called the “root zone distribution system” that is only used for distributing the root zone, the root-severs.net zone, and the .arpa zone to the RSOs. The RZM creates and distributes a new root zone approximately two or three times a day.

When the RZM prepares a root zone based on the data provided by IANA, it signs the zone using DNSSEC. The RZM uses the zone signing keys (ZSKs) that were signed by IANA during the key signing ceremonies (which are normally held four times a year). The RZM distributes the root zone with all additional DNSSEC resource records.

The RZM uses DNS NOTIFY messages defined in RFC 1996, “A Mechanism for Prompt Notification of Zone Changes (DNS NOTIFY)” to alert the RSOs that a new zone is available on the root zone distribution server. The DNS NOTIFY message is a method for a primary server to tell secondary servers that they should pull a new copy of the zone. In this case, the root zone distribution server acts as a “hidden primary” as defined in RFC 8499.

The RZM sends NOTIFY messages to a list of addresses provided by each of the RSOs. Each RSO can have multiple addresses on this list, based on each of the RSO’s internal mechanisms for distributing new copies of the root zones to the RSO’s instances. For example, one RSO might have each of its instances pull from the RZM’s server; another RSO might instead pull the zone into its management platform, and distribute the new root zone to its instances from there.

After receiving a NOTIFY message, the RSO transfers a new copy of the root zone from the root zone distribution server using AXFR (RFC 5936, “DNS Zone Transfer Protocol (AXFR)”). The zone transfers are currently authenticated with transaction signatures as described in RFC 2845, “Secret Key Transaction Authentication for DNS (TSIG).”

The RSOs react to the NOTIFY messages as they deem appropriate, and they use normal zone maintenance semantics when NOTIFY messages are not received in an expected timeframe. The RZM does not normally make any efforts to contact RSOs if they do not retrieve the new zone, but may do so in extraordinary circumstances such as emergency changes and during DNSSEC-related rollovers.

Operation of the root zone distribution server is described in the “Root Zone Maintainer Service Agreement” (often called the “RZMA”).

3 RSOs Serving the Root Zone

In order to serve the root zone correctly, every instance run by an RSO needs to have an up-to-date copy of the root zone. After an RSO has a copy of a new version of the root zone, it distributes that root zone to all of its instances. The method for pushing out the copies differs between RSOs. For example, some RSOs might use DNS NOTIFY messages and the DNS protocol's native zone transfer mechanism (AXFR) within their own set of servers, while others might use file transfer programs such as SFTP or rsync.

Each RSO responds to the DNS queries received using authoritative name server software running on each instance within the RSO's service. There are a variety of open source software packages for authoritative DNS service, and RSOs use many of these implementations. They can also use commercial authoritative name server software, software that they have written themselves, or open source software they have modified for their own purposes. Serving the DNS root zone often uses the same authoritative name server software used to serve other zones such as TLDs and registrants' zones.

RSOs often change the configuration of their authoritative software to maximize response rates. They also adjust the configuration of the operating system of the instances to better handle the expected load of UDP and TCP queries.

Each RSO decides how many instances it operates. The RSOs list the approximate number and location of their instances on a collective website managed by the RSOs. The number of instances for a particular RSO changes over time as instances are added and removed. Sometimes an RSO will add and remove an instance in a short period of time, which makes it appear that an instance moved. While deploying instances is beneficial to the resilience and performance of root service (which is a technical consideration), there are often non-technical factors leading to the deployment of instances including business, financial, and other considerations.

3.1 Operational Considerations for RSOs Serving the Root Zone

As described in RSSAC042, "RSSAC Statement on Root Server Operator Independence," each RSO makes its operational choices independently of the other RSOs. Each RSO makes many choices when operating their systems to serve the root zone. These include:

Number of instances – An RSO chooses how many instances it wants to support based on a wide variety of operational and business factors. An individual RSO might choose its instance count based on the geographic or business markets it wants to support, on the perceived reputation of its instance policy, and so on, even though the RSS as a whole is global and covers all business markets. When an RSO decides how many instances to deploy, it considers both its own desired operations strategy and the overall RSS. For example, an RSO that is considering adding one or more instances in a geographic area or a region of Internet topology, may also consider whether there are already instances from other RSOs that are adjacent or near.

Capabilities of instances – An RSO might choose to host only highly capable instances (those with good connectivity to the Internet and high DNS throughput), less capable instances (those whose connectivity and local bandwidth is more limited), or a mixture of instances across a spectrum of capabilities. An RSO may also consider the capabilities of instances operated by other RSOs that are topologically or physically close on the Internet. Another way that instances vary is in the number of computer systems that make up a single instance. Some RSOs configure many computer systems within a single instance in order to respond to a greater number of DNS queries.

Choice of authoritative server software – There are many choices of authoritative name server software with different features, such as support for some newer DNS features, different configuration capabilities, and amount of logging. All RSOs are committed to choosing name server software that correctly implements IETF standards. An RSO can also develop and deploy custom authoritative name server software.

Number of types of authoritative name server software to run – Running a single type of software on all instances leads to easier software maintenance. Running more than one type of software can give better overall resilience for the RSO if one of the types of software has a catastrophic vulnerability.

Configuration of each type of authoritative name server software – Typical server configuration choices include, among others, the level of logging of queries, and which DNS protocol options to support beyond those that are required for root service.

Configuration of operating systems – Tuning the operating system upon which a root server is running can cause major changes in the performance of the server. Configuring the operating system's TCP parameters related to the connection set-up and tear-down can be particularly important.

Routing operations – Universally, routing to instances relies on "anycast" technology, as described in RFC 4786, "Operation of Anycast Services." Thus, RSOs are heavily dependent on the routing configuration that is used in locations where the instances reside. Also, like routing for any non-trivial network, the routing needs to be maintained and changed as the routing in the nearby networks' topology changes, as peering partners shift, transit service is added and removed, and so on.

Local network configuration – Instances often exist in local networks that involve multiple servers and other devices. RSOs may need to tune these local networks to work with the types of traffic typically seen at instances, particularly for network options such as the maximum transmission unit (MTU).

Monitoring and reporting – RSOs are expected to monitor the health of their instances and fix any issues that appear in order to reduce the chance that a querier tries to use an instance that is not working. In RSSAC002v4, "Advisory on Measurements of the Root Server System," RSSAC recommends that each RSO reports various technical metrics to help monitor service trends in the RSS.

4 Resolvers Querying the Root Servers

The reason the RSS exists is for recursive resolvers to be able to start the process of DNS resolution to answer queries they receive. Root server instances do this by telling resolvers how to find the authoritative name servers for TLDs in the query names that the instances receive. There are millions (possibly tens of millions) of caching recursive resolvers that send queries to the root servers. The operation of these resolvers is important to the overall RSS because it affects the load on each of the root servers.

When a resolver starts up, it first gets the definitive list of root servers through a process called “priming," which is described in RFC 8109, “Initializing a DNS Resolver with Priming Queries.” The resolver then chooses one or more root servers to send all queries that it cannot answer from its cache. Many resolvers use one authoritative name server for the root at a time, and they often choose which authoritative name servers to use based on which responded most quickly to its priming queries. Note that resolvers treat root servers like other authoritative servers in this regard. Resolvers periodically re-check which authoritative name server to use because sometimes different servers become faster at responding.

Because recursive resolvers cache the responses from the root servers and all other authoritative name servers, queries received by the recursive resolver from stub resolvers can often be answered from the resolver’s cache instead of sending the query to a root server. Resolvers also have a “negative cache”: a cache of names for which there was a negative answer from the associated authoritative name server. Some resolvers use a technique called “aggressive negative caching” to greatly reduce the number of queries sent to the root servers and other authoritative servers.

4.1 Operational Considerations for Recursive Resolvers

The operator of a recursive resolver can often run the resolver without tuning its default configuration. However, there are many operational choices that resolver operators can make that affect the root server system. These include:

Choice of recursive resolver software – There are many choices of recursive resolver software with different capabilities, such as support for some newer DNS features that might affect authoritative name servers. Writing custom recursive resolver software is also an option, but one rarely taken due to the complexity of recursive resolvers.

Configuration of recursive resolver software – Additions and changes to the DNS protocol are made over time, and many of those affect recursive resolvers. Different implementations have different ways to configure whether those DNS features are turned on and how they are managed. For example, using the mechanism described in RFC 7816, “DNS Query Name Minimisation to Improve Privacy,” can reduce the size of queries going to the root servers.

DNSSEC validation –In order to perform DNSSEC validation, the DO (“DNSSEC OK”) bit in the EDNS0 extension must be set; some resolvers set this bit regardless of whether validation is being performed. Setting the DO bit causes responses from the root and other zones to be larger based on the presence of added DNSSEC metadata, and doing full DNSSEC validation can cause resolvers to send more queries to authoritative servers.

Using aggressive negative caching – A resolver that is performing DNSSEC validation can also use a mechanism, sometimes called “aggressive NSEC,” that uses the negative cache to reduce the number of queries sent to authoritative name servers for zones that are signed. This mechanism is described in RFC 8198, “Aggressive Use of DNSSEC-Validated Cache.” Because the root zone is signed using DNSSEC and uses NSEC records for proof of non-existence, a resolver that uses aggressive NSEC quickly reduces the number of queries sent to the authoritative name servers such as the root servers.

Memory allocation for the negative cache – Many of the queries sent to root servers are for names that do not exist in the DNS. If a resolver can be configured to have a larger cache, it will send fewer of these queries to the root servers.

Capping Time to Live (TTL) values – Some resolvers have configuration settings that allow them to cap the TTL of records going into their caches. Different resolver software has different default TTL caps. Limiting the TTL to two days or less will typically cause more queries from resolvers to be sent to root servers. This TTL cap applies to nameserver (NS) records for TLDs in the root zone as it does to all records that are cached.

Running a local root server – Resolvers that run a local root server, a process sometimes called “hyperlocal service,” will never send any queries to the root service under normal circumstances, as described in RFC 8806, “Running a Root Server Local to a Resolver.” These resolvers need to get up-to-date copies of the root zone from somewhere, and many RSOs allow resolvers to retrieve the root zone from dedicated servers.

5 Possible Future Changes to the RSS

In 2018, RSSAC published RSSAC037, “A Proposed Governance Model for the DNS Root Server System.” It is an extensive description of a model for adopting new governance structures to improve accountability and transparency in the RSS. It was accompanied by RSSAC038, “RSSAC Advisory on a Proposed Governance Model for the DNS Root Server System,” a very short document that proposes that the model in RSSAC037 be adopted by the ICANN Board.

Following the publication of those two documents, ICANN convened the Root Server System Governance Working Group (commonly called “the GWG”) to implement the ideas in RSSAC037. The GWG is looking at root server operations from the perspective of how the Internet community now supports the RSS, and how the community wants that to change in the future. It is also looking at how the community interacts with the RSS operators, and how that can evolve in the future. At the time of writing, the GWG is actively working on the model, and intends to present it to the ICANN Board for implementation.

At around the same time that the GWG was being formed, RSSAC published RSSAC047, “RSSAC Advisory on Metrics for the DNS Root Servers and the Root Server System.” The metrics in that document could be used by the new governance structure to hold RSOs to measurable minimum technical expectations. RSSAC047 also includes a set of proposed thresholds for the defined metrics.

5.1 Operational Considerations for Governance Evolution for the Root Server System

The new organization, as described in RSSAC037, would have the ability to potentially change the way that the current RSS is governed, based on what is likely to be extensive input from the community. The end result might be similar to today's RSS, but it also might be quite different. Some topics that the new organization may choose to consider include:

Target number of RSOs – The new organization might pick a target number of RSOs for the eventual RSS. Having fewer than 12 RSOs would simplify the management of the RSS, and could be done while keeping the same number of IP addresses because a single RSO can manage multiple addresses. Conversely, having more than 12 RSOs would likely increase the operational diversity of the RSOs.

Target number of instances – The new organization might pick a target number of instances for the eventual RSS. Having fewer than the current number of instances in the RSS would likely increase the latency seen by some recursive resolvers, but it is not known if this increase would be noticeable to the recursive resolver operators or Internet users. Having more than the current number of instances in the RSS would presumably increase the resilience of the RSS to DoS attacks because the attack traffic would be absorbed closer to the source.

Technical service targets – RSSAC047 lists four metrics and associated thresholds that might be applied to RSOs when evaluating their participation in the RSS in the future. The eventual list of metrics adopted by this new organization might be longer or shorter, and the thresholds given in RSSAC047 might be different.

Non-technical service targets – RSSAC001, "Advisory on Service Expectation of Root Servers," lists the expectations that all RSOs are required to meet. Some of these expectations are technical or operational and others are not. Section 3 of RSSAC037 also lists such principles. Expectations and principles that are hard to measure such as "good ethics" or "more diversity" might be used in the future for evaluating participation in the RSS.

6 Source Material

The following are the documents published by RSSAC that were used to prepare this document:

- RSSAC001, "Advisory on Service Expectation of Root Servers"
- RSSAC002v4, "Advisory on Measurements of the Root Server System"
- RSSAC023v2, "History of the Root Server System"
- RSSAC026v2, "RSSAC Lexicon"
- RSSAC037, "A Proposed Governance Model for the DNS Root Server System"
- RSSAC038, "RSSAC Advisory on a Proposed Governance Model for the DNS Root Server System"
- RSSAC042, "RSSAC Statement on Root Server Operator Independence"
- RSSAC047, "RSSAC Advisory on Metrics for the DNS Root Servers and the Root Server System"

Additional documents referred to in this document are:

- RFC 1996, “A Mechanism for Prompt Notification of Zone Changes (DNS NOTIFY)”
- RFC 2845, “Secret Key Transaction Authentication for DNS (TSIG)”
- RFC 4786, “Operation of Anycast Services”
- RFC 5936, “DNS Zone Transfer Protocol (AXFR)”
- RFC 7720, “DNS Root Name Service Protocol and Deployment Requirements”
- RFC 8109, “Initializing a DNS Resolver with Priming Queries”
- RFC 8198, “Aggressive Use of DNSSEC-Validated Cache”
- RFC 8806, “Running a Root Server Local to a Resolver”
- RFC 8499, “DNS Terminology”
- “Root Zone Maintainer Service Agreement”