# Key Technical Elements of Potential Root Operators

An Advisory from the ICANN Root Server System Advisory Committee (RSSAC)
October 14, 2016

# Preface

This is an Advisory to the Internet Corporation for Assigned Names and Numbers (ICANN) Board of Directors and the Internet community more broadly from the ICANN Root Server System Advisory Committee (RSSAC). In this Advisory, the RSSAC identifies key technical elements of potential DNS root server operators.

The RSSAC seeks to advise the ICANN community and Board on matters relating to the operation, administration, security and integrity of the Internet's Root Server System. This includes communicating on matters relating to the operation of the Root Servers and their multiple instances with the technical and ICANN community, gathering and articulating requirements to offer to those engaged in technical revisions of the protocols and best common practices related to the operational of DNS servers, engaging in ongoing threat assessment and risk analysis of the Root Server System and recommend any necessary audit activity to assess the current status of root servers and root zone. The RSSAC has no authority to regulate, enforce, or adjudicate. Those functions belong to others, and the advice offered here should be evaluated on its merits.

A list of the contributors to this Advisory, references to RSSAC Caucus members' statement of interest, and RSSAC members' objections to the findings or recommendations in this Report are at end of this document.

# Table of Contents

# 1. Introduction

In this document the RSSAC defines key technical elements of potential new root operators that would be a critical part of any potential root server operator designation process.

RSSAC001 ("Service Expectations of Root Servers") and RFC 7720 ("DNS Root Name Service Protocol and Deployment Requirements") are considered as starting points; alone, they are insufficient to evaluate potential operators.

Non-technical aspects, such as trustworthiness, ethos, funding, business models, openness, community participation and politics are out of scope for this document. The RSSAC believes these non-technical aspects to be important, and although this document does not address them, we expect them to be part of the overall evaluation. The elements defined in this document are designed to be:

- Technical in nature
- As specific as possible
- Provable, documentable, and/or measurable

The proposed recommendations only consider technical aspects as well as our current understanding of the key technical elements a potential root operator should meet. In other words, an organization meeting all elements in this document will not necessarily be designated as a root operator. Similarly, an organization that does not meet all elements might also be designated as a root server operator.

# 2. Assumptions

Since the root server designation process is yet-to-be-defined, this document makes certain assumptions for the sake of simplicity. It assumes that candidate operators will be provided a request for qualifications, to which each will provide a written response.

The technical aspects of the response will be evaluated based on the elements described in this document. This document provides no guidance on elements not listed in section 3 of this document.

This document represents RSSAC's input into the descriptions of technical elements for new root server operators. Other stakeholders might have additional or stronger requirements. A full procedure for evaluating candidate root server operators was not defined at the publication time of this document.

# 3. Technical Elements for Evaluation

This section contains the technical elements for a yet-to-be-defined evaluation process. Note that these are not written as requirements for operation, but rather elements for evaluation. Furthermore, this document does not necessarily describe *how* to evaluate those elements, only what they are.

## 3.1 RSSAC001 and RFC7720

The candidate operator must be evaluated with respect to existing operational expectations and requirements, namely RSSAC001 and RFC7720. Some of these might not be applicable to candidate operators. Others may be evaluated against a candidate's already existing services. Summaries of these expectations and requirements are given in Appendices A and B.

## 3.2 Design

### 3.2.1 Overall Service Design

The candidate operator's overall design must be evaluated with respect to its utility in serving the root zone. The candidate operator should provide as many design details as possible. Design choices might include hardware platforms, networking technology, use of virtualization, locations of servers (e.g., data centers, exchange points, shared cabinets), overall capacity, and out-of-band access.

### 3.2.2 Service Availability

The candidate operator's proposal must be evaluated with respect to its approach to maximizing service availability. Per RSSAC001 E.3.3-B, the candidate operator's design is expected to eliminate or minimize single points of failure. This might include diversity elements described in section 3.4.

### 3.2.3 Service Capacity

The candidate operator's service capacity must be evaluated for its ability to withstand Denial of Service (DoS) and other forms of attacks. See RSSAC001 E.3.4-A.

### 3.2.4 Performance

The candidate's design should be evaluated with respect to its performance characteristics, such as latency, serviced regions, RSSAC002 metrics and RFC7720 requirements.

## 3.3 Experience and Networking

### 3.3.1 DNS Operational Experience

Previous or current experience operating large-scale DNS services by the candidate operator should be considered. Such operation is expected to include both IPv4 and IPv6, and both UDP and TCP. Existing services of the candidate operator should be evaluated for similarities to the candidates expected root server operation (e.g., query rate, use of anycast, zone size, zone update frequency).

### 3.3.2 Security Audit

A security audit of the candidate operator should be performed and will be evaluated with respect to best current practices. It is expected that any security audit will be conducted by an organizations unaffiliated with the candidate operator. Audit results must be kept private unless otherwise agreed to by candidate operator.

### 3.3.3 Addressing Resources

The candidate operator must obtain its own AS number(s) and IPv4 and IPv6 address allocations for operating a root server. It is assumed that IP anycast will be used. If IP anycast will not be used, a technology providing similar or better service levels should be specified. Provider address space or addresses that cannot be used with anycast are undesirable. The expected production IPv4 and IPv6 address blocks must be severable from the candidate's organization to facilitate emergency or planned transfers.

### 3.3.4 DNS PTR Records

The candidate operator should demonstrate the ability to set DNS PTR records for its IPv4 and IPv6 address space.

### 3.3.5 Address Reputation

The reputation of the candidate operator's IP address blocks should be evaluated. Addresses with a bad reputation that are listed in one or more black lists (e.g., Spamhaus Don't Route Or Peer List) might affect a client's ability to reach the candidate's servers.

### 3.3.6 Peering Data

The candidate operator should have accurate and up-to-date information in known routing databases. If peering is to be used, accurate and up-to-date entries in a known peering database (e.g., PeeringDB) with complete information (e.g., contacts, policies, peers) are desirable.

### 3.3.7 Address Registries

The candidate operator's address space should be accurately registered in one of the Regional Internet Registry (RIR) public databases. Additionally, the candidate should have appropriate entries in relevant public routing registries for their IPv4 and IPv6 address space.

### 3.3.8 Zone Distribution Architecture

For efficiency, the candidate operator must maintain an internal zone distribution system. Since it is anticipated that the candidate operator will utilize IP anycast, load balancing, and/or multiple backend servers, the operator is expected to not unduly burden the root zone maintainer with an excessive number of zone transfer clients. The candidate operator should describe, in detail, their existing or proposed internal zone distribution architecture.

## 3.4 Diversity

Diversity *within* a given operator is of benefit, and in some instances diversity *among* operators may also be of benefit.

### 3.4.1 Geographic Diversity

The candidate operator is expected to provide root zone service from multiple geographic locations utilizing IP anycast, or a technology affording the same functionality as IP anycast. Serving more locations regionally will be considered better than serving less. The ability to operate in multiple continents and countries is preferred. Ideally, the candidate operator will provide service in regions not already well-served by other root operators.

### 3.4.2 Network Provider Diversity

Utilizing multiple upstream network providers can be of benefit to candidate operators that use third-party network providers. Candidate operators should demonstrate or document that they are not susceptible to the problems and sustained outages of a single network provider.

### 3.4.3 Network Hardware Diversity

Platform and vendor diversity can improve resilience by not relying on a single vendor or model for routers, switches, load balancers, and other networking equipment. For example, if a "0-day" vulnerability exists in a certain vendor's platform, the candidate operator should be able to continue using other, unaffected equipment.

### 3.4.4 Server Diversity

This can refer to different hardware vendors or different models of general purpose computers from the same vendor. Server diversity can be of benefit, whether it is inter- or intra-site.

### 3.4.5 Operating System Diversity

Operating system diversity allows an operator to continue operation in the event of an operating system defect causing an outage. Per RSSAC001 E.3.6-A, systemic operating system diversity is desirable.

### 3.4.6 Application Diversity

Application diversity provides resilience at the application layer, which includes name server software (e.g., BIND, Knot, NSD), routing software (e.g., Quagga, OpenBGPD, BIRD), and other possibly required applications. Per RSSAC001 E.3.6-A, systemic application diversity is desirable.

### 3.4.7 Human Diversity

The candidate operator should demonstrate or document that it does not rely on any single individual for technical operations. It is very beneficial when skills and knowledge are distributed so that operation continues even if key personnel depart or otherwise become unavailable.

### 3.4.8 Access Segregation

The candidate operator should segment which staff (e.g., engineering personnel, NOC staff, remote hands) have access to which equipment (e.g., servers, routers, load balancers). Documentation of which staff members have access to which equipment must be considered private to the candidate operator.

## 3.5 Documentation

### 3.5.1 Maintenance Procedures

The candidate operator should have documented maintenance procedures and make them available for evaluation. Per RSSAC001 E.3.3-A, the candidate operator should have the ability to take a subset of their service offline for maintenance without affecting the overall operation.

### 3.5.2 Emergency and Attack Recovery

In the event of unplanned outages, documented procedures should describe how to recover. For example, when anycast is utilized, routes are expected to be withdrawn from sites unable to provide service, either manually or automatically. The candidate operator should have documentation describing remote access, and how support staff can interact with a hosting provider's onsite support.

### 3.5.3 Disaster Recovery and Business Continuity

The candidate operator should have disaster recovery and business continuity plans. This includes recovering from natural disasters and other catastrophic events. This documentation should include information on backup sites, data recovery, and backup Network Operation Centers (NOCs).

### 3.5.4 Network Operation Center

The candidate operator should document their Network Operation Center (NOC). Documentation should include NOC availability, number of staff, time to respond, and staff on-call policies.

### 3.5.5 Computer Emergency Response Team Interaction

The candidate operator should have relationships and documented procedures in place for interacting with the larger security community and security advisories. This includes both responding to security advisories that might affect candidate operator's service, and advising the Internet security community of issues or attacks discovered or experienced by the candidate operator. Established relationships with local Computer Emergency Response Teams (CERTs) or Computer Security Incident Response Teams (CSIRTs) are advantageous.

## 3.6 Miscellaneous

### 3.6.1 Data & Measurement

The candidate operator should provide sample data collection output, including but not limited to RSSAC002 metrics. The candidate should commit to participate in Domain Name System-Operations Analysis and Research Center (DNS-OARC)[1] and its regular Day in the Life of the Internet (DITL)[2] data collections. If necessary, the candidate operator could be provided sample input data to test their data collection output. See RSSAC001 E.3.7-B.

### 3.6.2 Sample "x.root-servers.org" Web Page

The candidate operator should demonstrate their ability to maintain a <LETTER>.root-servers.org web page by providing a mock-up in HTML.

### 3.6.3 Evaluation Period

The candidate operator may undergo an evaluation period, the purpose of which is to demonstrate an ability to meet certain requirements specified in this document. During this time, the candidate's server is not published in the root zone nor the root-servers.net zone. Arrangements could be made to send query traffic to the candidate operator's servers using load-generation tools during this evaluation period.

[1] See https://www.dns-oarc.net/
[2] See https://www.dns-oarc.net/oarc/data/ditl

# 4. Acknowledgments, Disclosures of Interest, Dissents, and Withdrawals

In the interest of transparency, these sections provide the reader with information about four aspects of the RSSAC process. The Acknowledgments section lists the RSSAC caucus members, outside experts, and ICANN staff who contributed directly to this particular document. The Statement of Interest section points to the biographies of all RSSAC caucus members. The Dissents section provides a place for individual members to describe any disagreement that they may have with the content of this document or the process for preparing it. The Withdrawals section identifies individual members who have recused themselves from discussion of the topic with which this Advisory is concerned. Except for members listed in the Dissents and Withdrawals sections, this document has the consensus approval of the RSSAC.

## 4.1 Acknowledgments

RSSAC thanks the following members of the Caucus and external experts for their time, contributions, and review in producing this Report.

**RSSAC Caucus members**

Wes Hardaker
Howard Kash
Warren Kumari
Daniel Migault
Russ Mundy
Duane Wessels
Kevin Wright

**ICANN support staff**

Andrew McConachie (editor)
Steve Sheng

## 4.2 Statements of Interest

RSSAC caucus member biographical information and Statements of Interests are available at:
https://community.icann.org/display/RSI/RSSAC+Caucus+Statements+of+Interest.

## 4.3 Dissents

There were no dissents.

### 4.4 Withdrawals

There were no withdrawals.

# 5. Appendices

## Appendix A: Summary of Expectations from RSSAC001

[E.3.1-A] Individual Root Server Operators are to publish or continue to publish operationally relevant details of their infrastructure, including service-delivery locations, addressing information and routing (e.g., origin autonomous system) information.

[E.3.1-B] Individual Root Servers will deliver the service in conformance to IETF standards and requirements as described in RFC7720 and any other IETF standards-defined Internet Protocol as deemed appropriate.

[E.3.2-A] Individual Root Servers will adopt or continue to implement the current DNS protocol and associated best practices through appropriate software and infrastructure choices.

[E.3.2-B] Individual Root Servers will serve accurate and current revisions of the root zone.

[E.3.2-C] Individual Root Servers will continue to provide “loosely coherent” service across their infrastructure.

[E.3.2-D] All Root Servers will continue to serve precise, accurate zones as distributed from the Root Zone Maintainer.

[E.3.3-A] Individual Root Servers are to be deployed such that planned maintenance on individual infrastructure elements is possible without any measurable loss of service availability.

[E.3.3-B] Infrastructure used to deploy individual Root Servers is to be significantly redundant, such that unplanned failures in individual components do not cause the corresponding service to become generally unavailable to the Internet.

[E.3.3-C] Each root server operator shall publish documentation that describes the operator’s commitment to service availability through maintenance scheduling and notification of relevant operational events.

[E.3.4-A] Individual Root Server Operators will make all reasonable efforts to ensure that sufficient capacity exists in their deployed infrastructure to allow for substantial flash crowds or denial of service (DoS) attacks.

[E.3.4-B] Each Root Server Operator shall publish documentation on the capacity of their infrastructure, including details of current steady-state load and the maximum estimated capacity available.

[E.3.5-A] Individual Root Server Operators will adopt or continue to follow best practices with regard to operational security in the operation of their infrastructure.

[E.3.5-B] Root Server Operators shall publish high-level business continuity plans with respect to their Root Server infrastructure.

[E.3.6-A] Each Root Server Operator shall publish documentation that describes key implementation choices (such as the type of DNS software used) to allow interested members of the Internet community to assess the diversity of implementation choices across the system as a whole.

[E.3.7-A] Each Root Server Operator will adopt or continue to follow best current practices with respect to operational monitoring of elements within their infrastructure.

[E.3.7-B] Each Root Server Operator will adopt or continue to perform measurements of query traffic received and shall publish statistics based on those measurements.

[E.3.8.1-A] Individual Root Server Operators will continue to maintain functional communication channels between each other in order to facilitate coordination and maintain functional working relationships between technical staff.

[E.3.8.1-B] All communications channels are to be tested regularly.

[E.3.8.2-A] Individual Root Server Operators shall publish administrative and operational contact information to allow users and other interested parties to escalate technical service concerns.

# Appendix B: Summary of Requirements from RFC7720

### Protocol Requirements

The root name service:

- MUST implement core DNS [RFC 1035] and clarifications to the DNS [RFC 2181].

- MUST support IPv4 [RFC 791] and IPv6 [RFC 2460] transport of DNS queries and responses.

- MUST support UDP [RFC 768] and TCP [RFC 793] transport of DNS queries and responses.

- MUST generate checksums when sending UDP datagrams and MUST verify checksums when receiving UDP datagrams containing a non-zero checksum.

- MUST implement DNSSEC [RFC 4035] as an authoritative name service.

- MUST implement extension mechanisms for DNS (EDNS(0)) [RFC 6891].

### Deployment Requirements

The root name service:

- MUST answer queries from any entity conforming to [RFC 1122] with a valid IP address.

- MUST serve the unique [RFC 2826] root zone [ROOTZONE].