

Service Expectations of Root Servers

RSSAC-001¹, 2013-05-02

Table of Contents

Revision History	2
1. Introduction	2
2. Service Provided by Root Servers	3
3. Expectations of Root Server Operators	4
3.1. Infrastructure	4
3.2. Service Accuracy	4
3.3. Service Availability	5
3.4. Service Capacity	5
3.5. Operational Security	6
3.6. Diversity of Implementation	6
3.7. Monitoring and Measurement	6
3.8. Communication	6
3.8.1. Inter-Operator Communication	6
3.8.2. Public Communication	7
4. Public Documentation	7
5. Document Authors	7
6. Bibliography	8
Appendix A: Summary of Expectations	9

¹ Document number to be assigned, suggest 001

² All root servers apart from J-Root currently serve the ARPA zone.

³ For example, an IPv6 addresses was added to D-Root on 2011-06-10, and to I-Root

Revision History

This section to be removed upon final publication.

Revision	Date	Change
0.2	2012-05-09	Received baton from Terry.
0.3	2012-05-09	Extensive battle with Microsoft Word. Unclear as to who won. Also incorporated text to address comments from Matt, Joao.
0.4	2012-06-01	Stole baton from Joe, rev'd doc.
0.5	2012-06-27	Accepted changes by Matt
0.6	2013-04-22	Accepted Changes from Kato
0.7	2013-05-02	Added edits from John B. Completed appen A.
0.8	2013-08-07	Final suggested edits applied.

1. Introduction

Domain Name System (DNS) infrastructure includes elements known as Root Name Servers (“Root Servers”). This document describes the best practice service provided by Root Servers, and defines the expectations that users might reasonably hold of both that service and the Root Server Operators.

This document recognizes earlier guidance in the implementation and operation of Root Servers (RFC 2010) (RFC 2870), and the part such guidance has played in the development of the DNS as a whole. Earlier guidance provided detailed requirements on the technical implementation of root name servers that was useful at the time it was written. However, technical approaches for deploying authoritative-only DNS servers have advanced since that time, and there is a useful diversity of implementation evident in the root server system as a whole today that would not be possible if the strict advice in earlier documents were to be followed precisely.

This document highlights that a diversity of approach is desirable in the root server system, and replaces earlier direction on implementation with a set of service expectations that root server operators must satisfy.

It is left to the IETF to update and reclassify the earlier guidance in RFC28670 as historic.

2. Service Provided by Root Servers

At the time of writing there are thirteen Root Servers, operated by twelve different organizations. Root Servers are named A.ROOT-SERVERS.NET through M.ROOT-SERVERS.NET, and are often referred to by letter (e.g. "L-Root").

Although the word "server" is still used to identify the infrastructure providing service for individual letters, service is generally provided using techniques that involve more elaborate infrastructure than is suggested by that word. For example, many Root Servers provide service using multiple individual name server elements using anycast (BCP 126, RFC 4786), rather than being provided by a single server. In this document, "Root Server" refers generally to the service provided by the infrastructure operated by a Root Server Operator, and not to individual infrastructure elements.

From a protocol perspective, a Root Server is a DNS name server that provides authoritative-only DNS service for the root zone (STD 13, RFC 1034). Such name servers receive queries from clients using the DNS protocol (STD 13, RFC 1035) and provide appropriate responses. Clients of Root Servers are, for the most part, caching DNS resolvers that send requests to authoritative-only servers in response to queries they receive from stub resolvers.

Root Servers also serve additional zones. All Root Servers are authoritative for the ROOT-SERVERS.NET zone and currently twelve of the thirteen are authoritative for the ARPA zone².

The root zone of the DNS was signed using DNS Security Extensions (DNSSEC) (RFC 4033) in July 2010. Root Servers support the corresponding DNS protocol extensions (RFC 4035) when sending responses.

Each Root Server listens for queries on a set of IP addresses that are globally unique, and that are dedicated for use by that Root Server. At the time of writing some Root Servers listen on a single IPv4 address, and some listen on both a single IPv4 address and a single IPv6 address. Root Servers are renumbered occasionally, although such events are not frequent.³ Changes in service addresses for Root

² All root servers apart from J-Root currently serve the ARPA zone.

³ For example, an IPv6 addresses was added to D-Root on 2011-06-10, and to I-Root on 2010-06-17. F-Root's IPv6 address was renumbered on 2008-01-22. A summary of historical addressing changes can be found at <http://www.root-servers.org/>

Servers are coordinated by the IANA function⁴ as part of the normal Root Zone Management process.

3. Expectations of Root Server Operators

3.1. Infrastructure

[E.3.1-A] Individual Root Server Operators are to publish or continue to publish operationally relevant details of their infrastructure⁵, including service-delivery locations, addressing information and routing (e.g., origin autonomous system) information.

The public availability of this technical information facilitates troubleshooting and general operational awareness of Root Server infrastructure by the Internet technical community. The granularity of this information is limited to the publicly exposed service and at the comfort level of the Root Server Operator.

[E.3.1-B] Individual Root Servers will adopt or continue to deliver TCP and UDP DNS service over IPv4 and IPv6 and any other IETF standards-defined Internet Protocol as deemed appropriate.

3.2. Service Accuracy

[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.

The root zone content changes regularly although the extent of individual changes is generally small. Note, however, that at the time of this writing, the entire root zone is currently resigned every time it is published, so the DNSSEC signatures (i.e., RRSIG records) change with each new zone.

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

A set of name servers serving a single zone is said to be “loosely coherent” since although (ordinarily) all name servers in the set serve the same revision of the zone, there will be short intervals following the initial publication of a new revision of the zone in which some servers are observed to serve the now former zone, whilst others serve the newly published zone. The service provided by all 13 root servers by collective inheritance is similarly loosely coherent. Even though this ‘loosely

⁴ Internet Assigned Numbers Authority, <http://www.iana.org/>

⁵ A summary of all information published by root server operators, as described in this document, can be found in a later section.

coherent' paradigm exists, Root Server Operators will not impose any artificial delays on publishing a new revision of the Root Zone.

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

No Root Server has ever, or will ever, serve a zone that was modified following distribution by the Root Zone Maintainer. In any case, it would be impossible for an individual operator to modify the zone now that it is DNSSEC-signed without invalidating signatures. A Root Server Operator will not intentionally serve an older zone than current zone provided by the Root Zone Maintainer.

3.3. Service Availability

[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.

That is, there ought to be no planned maintenance associated with the operation of any Root Server that would make the corresponding service generally unavailable to the Internet.

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

There has been no documented example of a simultaneous failure of all Root Servers. The DNS protocol accommodates unavailability of individual Root Servers without significant disruption to the DNS service experienced by end users.

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

3.4. Service Capacity

[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.

Such events might present a significantly greater query load than the observed steady state, and that abnormal load should be accommodated, where possible and within reason, without degradation of service to legitimate DNS clients. Filtering techniques may be employed by Root Server Operators to maintain service to legitimate DNS queries.

[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.

Service Expectations of Root Servers

A root server operator might choose to publish its maximum estimated capacity in high-level terms to avoid disclosing operationally sensitive information that would potentially serve to provoke attackers.

3.5. Operational Security

[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.

This provides confirmation to the Internet community that disaster recovery plans exist and are regularly reviewed and exercised.

3.6. Diversity of Implementation

[E.3.6-C] 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.

Individual Root Server Operators make implementation decisions autonomously, but in a coordinated fashion. In particular, Root Server Operators collaborate to ensure that a diversity of software and related service-delivery platform choices exists across the Root Server system as a whole. The goal of this diversity is to ensure that the system as a whole is not unnecessarily dependent on a single implementation choice, which might otherwise lead to a failure of the whole system due to a serious defect in a common component.

3.7. Monitoring and Measurement

[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.

The goal here lies in identifying failures in service elements and mitigating those failures in a timely fashion.

[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.

The Internet technical community is then able to gauge trends and other effects related to production Root Server traffic levels.

3.8. Communication

3.8.1. Inter-Operator Communication

[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.

Emergency communications channels exist to facilitate information sharing between individual Root Server Operators in real time in the event that a crisis requires it.

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

3.8.2. Public Communication

[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.

4. Public Documentation

This document specifies that many aspects of the operation of individual Root Servers be published:

- Operationally relevant details of infrastructure, including service-delivery locations, addressing information and routing information.
- A commitment to service availability through maintenance scheduling and notification of relevant operational events.
- Infrastructure capacity, including details of current steady-state load and maximum estimated capacity available.
- High-level business continuity plans.
- Key implementation choices, such as the type of DNS software deployed.
- Statistics based on query traffic received.
- Operational contact information to allow escalation of technical service concerns.

All documentation is published at <http://www.root-servers.org/>⁶ or, if published elsewhere, is linked to from that page.

5. Document Authors

This draft document was prepared by the following individuals. It is anticipated that a future revision of this document, reviewed and reaching consensus within RSSAC, will be published as a collaborative product of that committee.

J. Abley, ICANN, L-Root

⁶ Information about root servers has long been published on <http://www.root-servers.org/>, a resource coordinated and managed collectively by all Root Server Operators.

Service Expectations of Root Servers

J. Damas, ISC, F-Root

M. Larson, Verisign, A- and J-Root

L-J. Liman, Netnod, I-Root

T. Manderson, ICANN, L-Root

B. Verd, Verisign, A- and J-Root

6. Bibliography

- Abley, J., & Lindqvist, K. (2006). *BCP 126, RFC 4786*. Operation of Anycast Services.
- Abley, J., Manderson, T., Damas, J., Verd, B., & Larson, M. *draft-rssac-rfc2010-rfc2870-to-historic*. Reclassification of RFC 2010 and RFC 2870 as Historic.
- Arends, R., Austein, R., Larson, M., Massey, D., & Rose, S. (2005). *RFC 4033*. DNS Security Introduction and Requirements.
- Arends, R., Austein, R., Larson, M., Massey, D., & Rose, S. (2005). *RFC 4035*. Protocol Modifications for the DNS Security Extensions.
- Bush, R., Karrenberg, D., Koster, M., & Plzak, R. (2000). *RFC 2870*. Root Name Server Operational Requirements.
- Manning, B., & Vixie, P. (1996). *RFC 2010*. Operational Criteria for Root Name Servers.
- Mockapetris, P. (1987). *STD 13, RFC 1034*. Domain names - concepts and facilities.
- Mockapetris, P. (1987). *STD 13, RFC 1035*. Domain names - implementation and specification.

Appendix A: Summary of Expectations

[E.3.1-A] Individual Root Server Operators are to publish or continue to publish operationally relevant details of their infrastructure⁷, including service-delivery locations, addressing information and routing (e.g., origin autonomous system) information.

[E.3.1-B] Individual Root Servers will adopt or continue to deliver TCP and UDP DNS service over IPv4 and IPv6 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 with no 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.

⁷ A summary of all information published by root server operators, as described in this document, can be found in a later section.

[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-C] 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.

The goal here lies in identifying failures in service elements and mitigating those failures in a timely fashion.

[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.