RSSAC028
Technical Analysis of the Naming Scheme Used For Individual Root Servers

An Advisory from the ICANN Root Server System Advisory Committee (RSSAC)
3 August 2017

Preface

This is a report 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 report, the RSSAC conducted a technical analysis of the naming scheme used for individual root servers.

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 those engaged in technical revisions of the protocols and best common practices related to the operation of DNS servers, engaging in ongoing threat assessment and risk analysis of the root server system and recommending any necessary audit activity to assess the current status of root servers and the 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 merit.

A list of the contributors to this report, references to RSSAC Caucus members' statements of interest and objections to the findings or recommendations in this report can be found near the end of this document.

Table of Contents

1. Introduction

The Domain Name System (DNS) is supported by root servers that serve the root zone. Individual root servers were named under the "root-servers.net" domain in 1995. The root-servers.net zone is delegated to the root servers.

This naming scheme has worked well for root servers and the Internet community at large for over two decades. However, given today's Internet environment, the RSSAC has studied the naming scheme used for individual root servers and considered the consequences of making changes.

The study documents a risk analysis of different alternative naming schemes. This analysis includes:

- Where the names reside in the DNS hierarchy
- Who administers the zone in which the names reside
- How different naming schemes affect DNSSEC validation of priming responses
- The size of priming responses[1]

From the risk analysis, the document aims at providing:

- Recommendation to root server operators, root zone management partners, and ICANN on whether changes should be made, and what those changes should be
- Recommendations on signing the addresses associated with the root servers
- Recommendation on the naming scheme for the root servers

1.1 Scope of Work

On July 9, 2015 the RSSAC issued a scope of work that provided direction for the work described in this document. As a courtesy to readers, the specified scope is included below, together with commentary on the treatment of each point provided in this document.

[1] While priming response size is of concern, we do not assume that smaller priming responses are necessarily better. The real concern is whether or not priming responses experience fragmentation.

RSSAC Scope of Work	Response
Document the technical history of the names assigned to individual root servers since the creation of the root server system.	*See RSSAC023, "History of the Root Server System"*
Consider changes to the current naming scheme, in particular whether the names assigned to individual root servers should be moved into the root zone from the ROOT-SERVERS.NET zone.	*See section 5*
Consider the impact on the priming response of including DNSSEC signatures over root server address records.	*See section 5 and Appendix A*
Perform a risk analysis.	*See section 6*
Make recommendations to root server operators, root zone management partners and ICANN on whether changes should be made, and what those changes should be.	*See section 7*

2. Terminology

In addition to the below terms, this document also uses common DNS terms from RFC 7719.[2]

Authoritative server – A system that responds to DNS queries with information about zones for which it has been configured to answer with the Authoritative Answer (AA) flag in the response header set to 1. It is a server that has authority over one or more DNS zones.

Delegation – A delegation is indicated by the presence of an NS RRset which associates a domain name to one or more server names. It indicates that the server names present in this RRset are authoritative for all labels below this domain name (unless there is a further delegation).

Glue records – Resource records within a response that are not part of authoritative data but are necessary in order to enable a resolver to complete the query resolution process under certain cases.[3]

In-zone – Records are in-zone for a server if that server is authoritative for those records.

[2] See RFC 7719, DNS Terminology. P. Hoffman, A. Sullivan, K. Fujiwara. December 2015.

[3] See Section 4.2.1 RFC 1034, Domain Names - Concepts and Facilities. P.V. Mockapetris. November 1987.

Key signing key (KSK) – A DNSSEC key that only signs the apex DNSKEY RRset in a zone. KSKs have the Secure Entry Point (SEP) flag set to 1.[4]

Node re-delegation attack – These attacks, if found to be feasible, could possibly allow an attacker to poison the cache of a recursive resolver in a similar fashion to the well-known "Kaminsky attack". Node re-delegation attacks[5] might affect the resolution of all zones in resolvers that do not validate, and all unsigned zones in validating resolvers. Section 7.2 recommends further study to determine whether these attacks are feasible and, if so, what the effects might be.

Priming resolution – The act of a resolver getting its initial set of addresses for the DNS root servers.[6]

Resolver – A program that retrieves information from name servers in response to client requests. A resolver performs queries for a name, type, and class, and receives answers. The logical function is called "resolution".[7]

Resource Record Set (RRset) – A set of resource records with the same label, class and type, but with different data.[8]

Zone signing key (ZSK) – A DNSSEC key that can be used to sign all the RRsets in a zone that require signatures, other than the apex DNSKEY RRset.[9]

3. Brief Functional Description of the Priming Process

The root servers are the authoritative servers for the root zone and are designated by a combination of NS and A/AAAA RRsets. The NS RRsets provide the domain names and the A/AAAA records provide the IP addresses for each record in the NS RRset.

In the priming resolution process, resolvers query for the NS RRset of the root; the response contains those NS records in the response's Answer section and some or all of the A/AAAA records in the response's Additional section.

In the current naming scheme, the responses for NS and A/AAAA records may or may not contain DNSSEC records, depending on whether a resolver requested them or not. Validating priming responses using DNSSEC enables a resolver to protect itself from

[4] See RFC6781, DNSSEC Operational Practices, Version 2. O. Kolkman, W. Mekking, R. Gieben. December 2012.

[5] See Improved DNS Spoofing Using Node Re-Delegation, https://www.sec-consult.com/fxdata/seccons/prod/downloads/whitepaper-dns-node-redelegation.pdf

[6] The reasons that a recursive resolver needs this information, and the mechanisms it can use to get it, are covered in *Initializing a DNS Resolver with Priming Queries* (IETF work in progress).

[7] See section 2.4 RFC 1034, Domain Names - Concepts and Facilities. P.V. Mockapetris. November 1987.

[8] See RFC 2181, Clarifications to the DNS Specification. R. Elz, R. Bush. July 1997.

[9] See RFC 6781, DNSSEC Operational Practices. Version 2. O. Kolkman, W. Mekking, R.Gieben. December 2012.

attacks that give incorrect addresses for the root servers. Currently, the root zone itself is signed, but the zone that authoritatively contains the root server addresses (that is, root-servers.net) is not. Therefore, responses in the priming resolution currently contain DNSSEC records for the NS but not the A/AAAA resource record sets.

4. Brief History of Names Assigned to Individual Root Servers

Please see RSSAC023 *History of the Root Server System*.

5. Analysis of Naming Schemes

This section describes various potential naming schemes for the root zone and associated root servers, including the current naming scheme. There are many characteristics that need to be considered when evaluating a naming scheme:

- Where the name resides in the DNS hierarchy
- Who administers the zone in which the names reside
- How the names can be validated with DNSSEC
- The size of priming responses

This section looks at different naming schemes, including:

1. The current naming scheme
2. The current naming scheme with the root-servers.net zone signed
3. In-zone NS names
4. Shared delegated TLD
5. Names delegated to each operator
6. Single shared label for all operators

Each of the schemes is further described from section 5.1 to section 5.6. Appendix A shows how recent authoritative servers would act for each of the scenarios given.

Other than the first scheme, all schemes intentionally have DNSSEC signatures over the addresses of the root zone’s nameservers either directly in the root zone (if no delegation occurs), or as a signed delegation.

In the schemes that use new short labels in the root zone, “a”, “b” and so on are used because those are the same names that are used today for the root servers. Further study might be needed to see if those short labels in the root zone will cause any significant problems.

Fragmentation may result in lost packets, either due to loss of fragments, or due to network equipment that blocks fragments. Resolvers should be able to recover from such losses by requesting smaller UDP sizes or by retrying over TCP. Individual root server operators may make different decisions on whether to allow fragmentation or to prevent it by specifying a limit on the size of UDP responses they will return.

Regardless of the scheme, the size of UDP responses is controlled by a negotiation between the resolver and the individual root server receiving the query. The size used will be the smaller of the configured value on the root server, and the size requested by the resolver. Depending on the scheme, smaller negotiated values may result in exclusion of RRSIGs, some or all glue addresses in the Additional section, or even truncation (with the response having TC=1 set), if the answer will not fit in the UDP response; such a response might cause the client to retry over TCP. The exclusion of RRSIGs or glue addresses may result in resolvers performing additional queries in order to obtain signatures.

5.1 The Current Naming Scheme

In the current naming scheme, the authoritative servers for the root zone have the names "[a-m].root-servers.net". The root-servers.net zone is served by name servers that also serve the root zone. In this scheme, the root-servers.net zone continues to be unsigned.

Advantages of the current scheme are:

- The zone split of the root-servers.net zone follows the traditional DNS rules and limits the risk of any misinterpretation.
- Maintains the status quo.

Drawbacks of the current scheme are:

- The root and root-servers.net zones need to be synchronized, and stay synchronized, because information associated with the root servers is located in the root zone, the .net zone, and in the .root-servers.net zone.
- Root servers are not authoritative for the .net zone. This means that if the authoritative servers for the .net zone were unavailable, it could prevent the resolution of the root-servers.net zone.
- Because the root server address records are not signed, there is the possibility for DNS-based spoofing attacks on the root server infrastructure.

5.2 The Current Naming Scheme, with DNSSEC

This is the same as the preceding scheme, but with root-servers.net being signed by the zone's maintainer.

Glue records are included in the root zone. However, because the root zone is not authoritative for these glue records, the root zone does not contain their associated RRSIG records; in this scheme, the root-servers.net zone would be signed. Different authoritative server software will act differently with respect to those glue records. Some authoritative server software will include the RRSIGs, others won't, depending on the configuration of the authoritative server software being used.

Possible advantages of this scheme are:

- The zone split used by the root-servers.net follows the traditional DNS best practices thereby limiting risk of any misinterpretation.

- Signing the root-servers.net zone enables DNSSEC-aware resolvers to cryptographically validate the priming response.

Possible drawbacks of this scheme are:

- The root, root-servers.net and net zones need to be synchronized, and stay synchronized, because information associated with the root servers is located in the root zone, the .net zone, and in the .root-servers.net zone.
- Root servers are not authoritative for the .net zone. This means that if the authoritative servers for the net zone were down, it could prevent the resolution of the root-servers.net zone.
- After the priming query, a validating recursive resolver must query the net zone for the NS and DS records, and then query root-servers.net in order to get the DNSSEC data. This results in additional round trips for the resolver.
- If the servers for the net zone were unavailable, the DS records for root-servers.net zone would not be obtainable and validation of the priming response would fail.

5.3 In-zone NS Names

The root zone will have an NS RRset consisting of in-zone names with the A and AAAA records of the root servers. Because the records are maintained in the root zone, there would be no delegation points and the root zone would be authoritative for all content required for a priming query response. In this proposal, the names can either have all records under a common undelegated subdomain (for example, the names “a.root-servers”, “b.root-servers”, and so on) or can be short labels in the root zone (for example, the names “a”, “b”, and so on).

Depending on the name server software and configuration, the response to a priming query would contain an Answer section with 13 NS records and an Additional section that may contain all 13 A and AAAA glue records and 26 RRSIG records.

Possible advantages of this scheme are:

- The names could be similar to the current lettering scheme.
- All data is protected by DNSSEC.
- The DNSSEC data could be returned in the first query. There is no DNSSEC chain to follow; and, in an ideal situation, all RRSIG records would be contained in the response.
- Authentication of priming query responses requires only the keys for one zone. There are no additional DS records or additional keys for subordinate zones.
- It is syntactically elegant because the zone is clearly authoritative for its own name servers. There is no ambiguity regarding where the content could be found.
- Administration is simplified because changes only require one entity, not a coordination between the maintainer of the root zone and the child zone.

Possible drawbacks of this scheme are:

- There may be name collisions from search lists[10] (similar to the possibility of name collisions that happen any time a new TLD is added to the root zone) for this new common undelegated domain or the short labels.
- The response size with a full additional section of a standard priming query would far exceed the common UDP packet sizes for both IPv4 (1500) and IPv6 (1280). There is reason to believe that many networks drop IPv6 extension headers (and thus may also drop fragmented IPv6 packets) as well as dropping ICMPv6 packets.

5.4 Shared Delegated TLD

The root zone will have an NS RRset that consists of 13 domain names that share a new common delegated TLD (for example, the names “a.root-servers”, “b.root-servers”, and so on). There will be 13 records in the root zone’s NS RRset pointing to the root server nameserver instances. The new shared TLD will be delegated to the same set of nameservers.

The response to a priming query has an Answer section with 13 NS records and an RRSIG for the NS RRset, and an Additional section with all the A and AAAA glue. Name server implementations differ in their behavior on whether the RRSIGs for these A and AAAA records are returned in the priming response. If the RRSIG RRset for the addresses is missing, a validating recursive resolver must query the root for the shared TLD’s NS RRset and then query the shared TLD for the A and AAAA RRsets.

It is possible to use an existing TLD that is hosted by the root servers, .arpa, for this proposal. However, that zone is administered by a different organization, the Internet Architecture Board (IAB), and thus using that TLD instead of a new one would mean that changes would need to be synchronized and approved outside the current set of involved actors.

Possible advantages of this scheme are:

- The names could be similar to the current lettering scheme.
- All data is protected by DNSSEC.
- The DNSSEC signatures all come from just one entity.
- Administration is simplified because changes only require one entity, not a coordination between the maintainer of the root zone and the child zone.

Possible drawbacks of this scheme are:

- The root and the common delegated TLD need to be synchronized, and stay synchronized, because information associated with the root servers is located in the common delegated TLD.
- As part of the priming query, a validating recursive resolver must query the root for the NS records, then query the shared TLD in order to get the DNSSEC data.
- There may be name collisions from search lists (similar to the possibility of name collisions that happen any time a new TLD is added to the root zone) for this new

[10] See SAC064 – SSAC Advisory on Search List Processing (13 February 2014)

shared TLD name.
- A new TLD would likely be subject to more policy, scrutiny and oversight than the current root-servers.net domain

5.5 Names Delegated to Each Operator

A new domain will be delegated to each root server operator. The root zone will have an NS RRset consisting of server names that are managed by the corresponding root server operators. The names for this proposal can either have all records under a common label (for example, the names “a.root-servers”, “b.root-servers”, and so on) or can be short labels in the root zone (for example, the names “a”, “b”, and so on). No other delegations are involved.

The response to a priming query has an Answer section with 13 NS records and an RRSIG for the NS RRset, and an Additional section with all the A and AAAA glue, but no RRSIG records for the address records. To get the RRSIG RRset, a validating recursive resolver must query the nameserver for each individual operator.

Possible advantages of this scheme are:

- The names could be similar to the current lettering scheme.
- All data could be protected by DNSSEC.
- The initial response might be small.

Possible drawbacks of this scheme are:

- After the priming query, a validating recursive resolver must query the root for the NS records for each operator’s TLD, then query the nameserver for each operator in order to get the DNSSEC data.
- There may be name collisions from search lists (similar to the possibility of name collisions that happen any time a new TLD is added to the root zone) for this new common domain or the short labels.
- Instead of just one entity signing a zone, each root zone operator needs to sign its own zone. This greatly increases the chances of operational error during the signing process, which may lead to some resolvers being unable to validate the priming response.
- Some root server operators might not sign their zone, or might want to sign with different algorithms from the other operators, which may result in other security or operational implications that have yet to be studied.

5.6 Single Shared Label for All Operators

Instead of having individual names for each root server, the set of root servers could be given one name at the top level (such as “all-root-servers.”) and that one name has the 13 IPv4 addresses and 13 IPv6 addresses of the root servers as two RRsets.

The response to a priming query has an Answer section with 1 NS record and an RRSIG, and an Additional section with all the A and AAAA glue and two RRSIG records (one each for the A and one for the AAAA RRsets).

Possible advantages of this scheme are:

- All addresses for the root servers are in only one place: the Answer section for the priming query.
- Administration is simplified as changes only require one entity, not a coordination between the maintainer of the root zone and the child zone.
- The DNSSEC data could be returned in a single query: There is no DNSSEC chain to follow, and in an ideal situation all RRSIG records would be contained in the response
- Validation of priming responses requires only the keys for only one zone. There are no additional DS records or additional keys for subordinate zones.
- It is syntactically elegant because the zone is clearly authoritative for its own name servers. There is no ambiguity regarding where the content could be found.

Possible drawbacks of this scheme are:

- If a resolver treats its configured addresses as RRsets instead of individual addresses, and if a priming query to any of the root servers results in a SERVFAIL or REFUSED response, resolvers might be unable to complete the priming query because they might not try to send queries to any of the other records in the A or AAAA RRset. Fixing this issue would require both protocol work and a full implementation rollout.
- There may be name collisions from search lists (similar to the possibility of name collisions that happen any time a new TLD is added to the root zone) for the single shared label.

6. Analysis of Benefits vs. Risks

The trade-offs between different naming alternatives are summarized in the table below.

Concerns	5.1 Current Naming Scheme	5.2 Current Naming Scheme with DNSSEC	5.3 In-zone NS RRset	5.4 Shared delegated TLD	5.5 Names delegated to each operator	5.6 Single Shared Label for all operators
Need to synchronize data in multiple zones	X	X		X	X	
External dependency on a zone not considered part of the root server infrastructure	X	X				
Exposure to DNS-based attacks on the root server infrastructure	X					
Maintaining the status quo	X					
Adding complexity due to DNSSEC signing		X	X	X	X	X
Increased workload associated with validating a longer authentication chain		X		X	X	
Increased round-trip delay associated with validating the priming response		X		X (some systems)	X	
Increase in priming response size		X	X	X	X	X
Corner cases and potential for errors					X	X
Name collision with search lists			X	X	X	X
Reduced root server operator autonomy						X

See Appendix A for a for a list of response sizes for each proposed scheme. Note that response sizes change depending on the type of authoritative software used and the configuration parameters chosen by the root server operator.

When the root-servers.net zone is unsigned, it exposes the DNS infrastructure to node re-delegation attacks on the addresses for the root servers. However, this naming scheme is known to work with the current resolver population. Maintaining the current status quo is an option if the risks associated with making changes to the root naming infrastructure outweigh the risks of re-delegation attack or expected benefits from a new naming scheme.

The risks associated with unsigned root server names can be mitigated by signing the zone that is authoritative for these names. A number of different naming schemes are possible here, and each scheme has its own unique set of concerns.

The option that is likely to involve the least change to the existing root server infrastructure is that of signing the root-servers.net zone (5.2). However, this approach brings with it the continued dependence on the .net zone, with the added burden of having to ensure that the secure delegation from net to root-servers.net remains valid.

The dependency on the .net zone can be removed by moving the root server names to the root zone (5.3) or to a new TLD under the root zone (5.4). There are trade-offs associated with each alternative. In the case of 5.3 the priming response size is largest on average. However, the additional information in the larger response also enables a validating resolver to authenticate the name server addresses without the need for additional lookups. In addition, since the root server names are authoritative data in the root zone, there is no secure delegation to follow while verifying the signatures covering these names.

In the case of 5.4 there is an additional overhead associated with managing and verifying the secure delegation from the root zone to the shared TLD. In option 5.4 the shared TLD and the root zone are both served by the root servers. However, different name server implementations differ on whether or not they return RRSIG information for the name server names within the shared TLD. In cases where these signatures are not returned there is an additional lookup overhead associated with fetching this information. In cases where these signatures are returned, the response size increases.

The advantage of option 5.4 is that it fails more gracefully if fragmented responses prove to be a problem. In the worst case, if a root server returns the aggregated information in the priming response there is little difference in the response sizes between 5.3 and 5.4. However, in cases where the polled root server's implementation does not include the complete set of A/AAAA information with signatures, fragmentation may not occur and clients may not see this breakage.

It is important to note that the potential of 5.4 to fail gracefully is only conjecture at this time. Additional studies are needed to verify this claim empirically.

Variant 5.5, where a separate delegation is made to each root server operator, may afford the root server operators greater flexibility and autonomy over the definition of the root server names. However, this flexibility comes at the cost of increasing the round-trip delay and overhead associated with signing and operating multiple signed zones, and for validating the A/AAAA RRsets for each root server operator managed zone. (There is currently experimentation with this scenario being performed by the Yeti DNS Project.)

In the final variant 5.6, consolidating all root servers under one name trades the overhead associated with managing multiple root server names for a larger A and AAAA RRset size. Because the number of RRSIGs covering the A/AAAA records is far fewer, this option also produces the smallest signed priming response that contains the full set of A and AAAA records associated with the root servers. However, this alternative may also result in new corner cases, such as in the way that query load is distributed across various root servers if resolvers identify different root servers through their names rather than their IP addresses.

All naming schemes that introduce a new TLD or a new name in the root zone increase the potential of name collisions with existing resolver search lists. Similarly, all naming schemes that involve a signed namespace for the root server names introduce a concomitant effect on the signed DNS response sizes. The level of size increase is different for the different options, as summarized in Appendix A.

7. Recommendations

7.1 Primary Recommendation

Recommendation 1: No changes should be made to the current naming scheme used in the root server system until more studies have been conducted.

Based on the investigation conducted by the RSSAC Caucus Root Server Naming Work Party, the near-term recommendation is that no changes should be made to the current root server system naming scheme. The work party concluded that there may be a benefit to later moving to one of the schemes listed in Section 5, based on the risk analysis explained in Section 6. However, it was recognised that more in-depth research is required to understand node re-delegation attacks, the costs and benefits of signing the A and AAAA records for the root servers, and the effects of increasing the priming query response size.

7.2 Further Studies

Recommendation 2: Conduct studies to understand the current behavior of DNS resolvers and how each naming scheme discussed in this document would affect these behaviours.

To better understand the findings of this report, DNS researchers should investigate the following topics, which have been covered earlier in this document. The operational differences between the options in Sections 5.3 and 5.4 are particularly relevant for further research. Some topics that would be of interest include:

- The acceptable response size (beyond the default UDP packet size) for priming queries. For example, IoT devices acting as DNS resolvers might not be able to receive long priming responses.
- How different resolver software responds when answers contain a reduced set of glue records.
- How current resolver implementations behave if they set the “DNSSEC OK” (DO) bit to 1 in their priming queries, such as if they validate the response and, if so, how they handle a bogus response.
- How search lists might be relevant. In the unusual case that a resolver also uses a DNS search list, using a single label for the root servers may interfere with that search list mechanism unless the final ‘.’ is given in the searched-for names.

If a change to the naming scheme is ultimately accepted, a transition plan would need to be produced to explore the practical obstacles faced by such a change. That transition plan itself would be a research topic.

Recommendation 3: Conduct a study to understand the feasibility and impact of node re-delegation attacks.

Further study is required to understand whether the current infrastructure is susceptible to various cache poisoning attack scenarios, including the cited node re-delegation attack. If the infrastructure is determined to be susceptible, the study needs to say what the effects of such attacks might be. Understanding these risks is necessary to assess the risk of changing the current root naming infrastructure. Any study conducted in this area should also be accompanied with proof-of-concept code so that it can be observed and further studied by the RSSAC Caucus and other researchers

Recommendation 4: Study reducing the priming response size.

When considering the priming response under DNSSEC, the scheme explained in Section 5.6 generated the smallest possible size, as expected. However, some implementations would become brittle if this naming scheme was adopted. Future work in this area could include modeling and proposing protocol changes to support this configuration, noting that the total cost shown by such a model might exceed the accompanying total benefit.

RSSAC should study having a specific upper limit on the size of priming responses where the query has DO=1. Research to reduce the response size might consider:

- Choosing a naming scheme with a single root server name
- Testing the consequences of all large responses having the TC bit set
- Backward-compatible protocol enhancements using EDNS0 to support a priming specific single signature over the entire priming set (NS, A, AAAA, DNSKEYs)

Further, more speculative studies about how to reduce the response size might include:

- Using different cryptographic algorithms
- Advertising what is expected in the Additional section (this would require modifying the DNS protocol)
- Having a single key for the root zone instead of the current KSK + ZSK scheme

- Effects of leaving the Additional section in priming responses empty

7.3 Speculative Recommendations

The fundamental recommendation of the RSSAC is to not change the current root server system naming scheme until the studies listed in section 7.2 can be completed. However, during the preparation of this document, the RSSAC Caucus Root Server Naming Work Party also made some observations that could be considered as recommendations based on particular outcomes in the further studies, and based on the risk analysis in Section 6.

If node re-delegation attacks pose a serious risk that needs to be mitigated, the following seem reasonable to consider:

- The root server addresses should be signed with DNSSEC to enable a resolver to authenticate resource records within the priming response. The root server addresses should be signed in a way that reduces the potential for operational breakage.
- Because the root server IP address information and the root zone are closely correlated, both sets of information should continue to be hosted on the same servers. This can be done using delegation or including the root server names in the root zone. All information necessary to validate the root-servers' A/AAAA RRsets and the root zone should be hosted on the root servers.
- Among the various options considered in this document, moving the root server names to the root zone (5.3), or adding a new TLD under the root zone (5.4) are both viable options that would result in signing the root server addresses. Additional studies are needed to determine which of these options, if any, would be more favorable than the other in practice.

8. 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.

8.1 Acknowledgments

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

RSSAC Caucus members
Joe Abley (work party leader)
John Bond (work party leader)
Brian Dickson
Paul Hoffman
Suresh Krishnaswamy
Warren Kumari
Matt Larson
Declan Ma
Bill Manning
Jim Martin
Robert Martin-Legene
Daniel Migault
Shinta Sato
Arturo Servin
Davey Song
William Sotomayor
Paul Vixie
Wesley Wang
Suzanne Woolf

ICANN Support Staff
Andrew McConachie (editor)
Kathy Schnitt
Steve Sheng (editor)

8.2 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

8.3 Dissents

There were no dissents.

8.4 Withdrawals

There were no withdrawals.

9. Revision History

9.1 Version 1

Current version.

Appendix A: Results from Testing Common Authoritative Servers

The test bed consists of recent versions of popular authoritative servers running with very minimal configurations.

The servers are running:

- BIND 9.10.3
- Knot 2.2.1
- Knot 2.3.0
- NSD 4.1.13

(Two recent versions of Knot were tested because there were significant technical changes between the two.) The zone files corresponding to the proposals were created by John Bond. The zone files can be AXFR'd from the addresses given in the configuration files.

Configuration Files

The configuration files used are listed on the following pages.

Proposal 5.1

```
knot.conf:
      server:
        listen: 0.0.0.0@53
        listen: ::@53
      remote:
        - id: master
          address: 2a03:b0c0:1:a1::189c:e001
      zone:
        - domain: "."
          master: master
        - domain: "root-servers.net"
          master: master
```

```
named.conf:
      options { recursion no; empty-zones-enable no ; dnssec-enable yes;
        listen-on { any; }; listen-on-v6 { any; }; };
      zone "." { type slave; masters {2a03:b0c0:1:a1::189c:e001;}; };
      zone "root-servers.net." { type slave; masters
{2a03:b0c0:1:a1::189c:e001;}; };
```

```
nsd.conf:
      zone:
        name: "."
        request-xfr: 2a03:b0c0:1:a1::189c:e001 NOKEY
        allow-notify: 2a03:b0c0:1:a1::189c:e001 NOKEY
      zone:
        name: "root-servers.net"
        request-xfr: 2a03:b0c0:1:a1::189c:e001 NOKEY
        allow-notify: 2a03:b0c0:1:a1::189c:e001 NOKEY
```

Proposal 5.2

```
knot.conf:
      server:
        listen: 0.0.0.0@53
        listen: ::@53
      remote:
        - id: master
          address: 2a03:b0c0:1:a1::189c:e002
      zone:
        - domain: "."
          master: master
        - domain: "root-servers.net"
          master: master
```

```
named.conf:
      options { recursion no; empty-zones-enable no ; dnssec-enable yes;
        listen-on { any; }; listen-on-v6 { any; }; };
      zone "." { type slave; masters {2a03:b0c0:1:a1::189c:e002;}; };
      zone "root-servers.net." { type slave; masters
{2a03:b0c0:1:a1::189c:e002;}; };
```

```
nsd.conf:
      zone:
        name: "."
        request-xfr: 2a03:b0c0:1:a1::189c:e002 NOKEY
        allow-notify: 2a03:b0c0:1:a1::189c:e002 NOKEY
      zone:
        name: "root-servers.net"
        request-xfr: 2a03:b0c0:1:a1::189c:e002 NOKEY
        allow-notify: 2a03:b0c0:1:a1::189c:e002 NOKEY
```

Proposal 5.3

```
knot.conf:
      server:
        listen: 0.0.0.0@53
        listen: ::@53
      remote:
        - id: master
          address: 2a03:b0c0:1:a1::189c:e003
      zone:
        - domain: "."
          master: master
```

```
named.conf:
      options { recursion no; empty-zones-enable no ; dnssec-enable yes;
        listen-on { any; }; listen-on-v6 { any; }; };
      zone "." { type slave; masters {2a03:b0c0:1:a1::189c:e003;}; };
```

```
nsd.conf:
      zone:
        name: "."
        request-xfr: 2a03:b0c0:1:a1::189c:e003 NOKEY
        allow-notify: 2a03:b0c0:1:a1::189c:e003 NOKEY
```

Proposal 5.4

```
knot.conf:
      server:
        listen: 0.0.0.0@53
        listen: ::@53
      remote:
        - id: master
          address: 2a03:b0c0:1:a1::189c:e004
      zone:
        - domain: "."
          master: master
        - domain: "root-servers"
          master: master
```

```
named.conf:
      options { recursion no; empty-zones-enable no ; dnssec-enable yes;
        listen-on { any; }; listen-on-v6 { any; }; };
      zone "." { type slave; masters {2a03:b0c0:1:a1::189c:e004;}; };
      zone "root-servers." { type slave; masters
{2a03:b0c0:1:a1::189c:e004;}; };
```

```
nsd.conf:
      zone:
        name: "."
        request-xfr: 2a03:b0c0:1:a1::189c:e004 NOKEY
        allow-notify: 2a03:b0c0:1:a1::189c:e004 NOKEY
      zone:
        name: "root-servers"
        request-xfr: 2a03:b0c0:1:a1::189c:e004 NOKEY
        allow-notify: 2a03:b0c0:1:a1::189c:e004 NOKEY
```

Proposal 5.5

```
knot.conf:
      server:
        listen: 0.0.0.0@53
        listen: ::@53
      remote:
        - id: master
          address: 2a03:b0c0:1:a1::189c:e005
      zone:
        - domain: "."
          master: master
        - domain: "a.root-servers"
          master: master
```

```
named.conf:
      options { recursion no; empty-zones-enable no ; dnssec-enable yes;
        listen-on { any; }; listen-on-v6 { any; }; };
      zone "." { type slave; masters {2a03:b0c0:1:a1::189c:e005;}; };
      zone "a.root-servers." { type slave; masters
{2a03:b0c0:1:a1::189c:e005;}; };
```

```
nsd.conf:
      zone:
        name: "."
        request-xfr: 2a03:b0c0:1:a1::189c:e005 NOKEY
        allow-notify: 2a03:b0c0:1:a1::189c:e005 NOKEY
      zone:
        name: "a.root-servers"
        request-xfr: 2a03:b0c0:1:a1::189c:e005 NOKEY
        allow-notify: 2a03:b0c0:1:a1::189c:e005 NOKEY
```

Proposal 5.6

```
knot.conf:
      server:
        listen: 0.0.0.0@53
        listen: ::@53
      remote:
        - id: master
          address: 2a03:b0c0:1:a1::189c:e006
      zone:
        - domain: "."
          master: master


named.conf:
      options { recursion no; empty-zones-enable no ; dnssec-enable yes;
        listen-on { any; }; listen-on-v6 { any; }; };
      zone "." { type slave; masters {2a03:b0c0:1:a1::189c:e006;}; };


nsd.conf:
      zone:
        name: "."
        request-xfr: 2a03:b0c0:1:a1::189c:e006 NOKEY
        allow-notify: 2a03:b0c0:1:a1::189c:e006 NOKEY
```

BIND 9.10.3

	No EDNS	IPv4 No DNSSEC MTU=16384	IPv4 DNSSEC MTU=16384	IPv6 DNSSEC MTU=16384
5.1	508	811	1097	1097
5.2	508	811	3833	3833
5.3	507	782	3938	3938
5.4	504	807	4093	4093
5.5	264	275	561	561
5.6	250	625	1485	1485

NSD 4.1.13

	No EDNS	IPv4 No DNSSEC MTU=16384	IPv4 DNSSEC MTU=16384	IPv6 DNSSEC MTU=16384
5.1	5.1	492	811	1097
5.2	5.2	492	811	1097
5.3	5.3	491	782	1418
5.4	5.4	488	807	1093
5.5	5.5	500	847	1133
5.6	5.6	250	625	1485

Knot 2.2.1

	No EDNS	IPv4 No DNSSEC MTU=16384	IPv4 DNSSEC MTU=16384	IPv6 DNSSEC MTU=16384
5.1	508	811	1097	1097
5.2	508	811	1097	1097
5.3	507	782	3938	3938
5.4	504	807	1093	1093
5.5	500	847	1133	1133
5.6	250	625	1485	1485

Knot 2.3.0

	No EDNS	IPv4 No DNSSEC MTU=16384	IPv4 DNSSEC MTU=16384	IPv6 DNSSEC MTU=16384
5.1	228	239	525	525
5.2	228	239	525	525
5.3	507	782	3938	3938
5.4	224	235	521	521
5.5	264	275	561	561
5.6	250	625	1485	1485