

Signing the root-servers.net zone

Summary

The root-servers.net zone was created in 1995 as a new way to name the root name servers. Since then, all the root name servers have been given single-letter names in this zone, such as *k.root-servers.net*. The primary purpose of the root-servers.net zone is to be the authoritative location for root server IP addresses.

Although the root zone was signed with DNSSEC in 2010, the root-servers.net zone was not. This was an intentional decision by the Root DNSSEC Design Team (true?) At the time, it was felt that signing root-servers.net was an extra complication, and not strictly necessary due to the way that DNSSEC is designed to protect DNS “leaf” data – i.e. data requested by end users.

In 2017, the RSSAC published its analysis of the naming scheme for individual root servers,¹ much of which focused on having signed data for root server names and addresses. Although the report does not recommend any changes to the current naming scheme, it does discuss the consequences of signing the root-servers.net zone.

We ask RZERC to consider the tradeoffs and provide recommendations regarding the DNSSEC status of the root-servers.net zone.

Proposal Details

In the design of DNSSEC, only authoritative zone data is signed. Non-authoritative data, and glue data in particular, is not signed. In the root zone, the A and AAAA records conveying the root server IP addresses are non-authoritative and are not signed. This was an intentional design decision. It means that DNSSEC can tell you whether or not you got the correct data, but not whether or not you got it from the correct server. In other words, DNSSEC doesn't care *where* data comes *from*, only whether or not it has been modified.

Figure 1 shows the relationship between the root zone, the root-servers.net zone, and which types of data are signed, not signed, and could be signed.

¹ RSSAC028 “Technical Analysis of the Naming Scheme Used For Individual Root Servers,” <https://www.icann.org/en/system/files/files/rssac-028-03aug17-en.pdf>

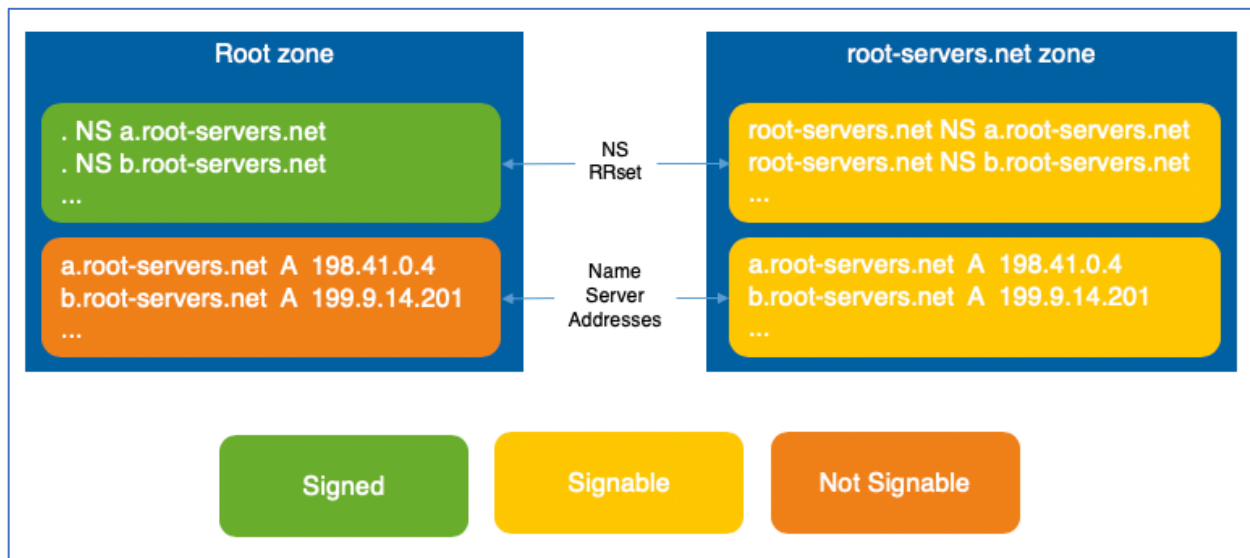


Figure 1. DNSSEC status of data in the root and root-servers.net zones

Signing the root-servers.net zone is of benefit to validating recursive name servers (aka “validators”). The clear advantage being the ability to ensure that root zone queries go only to real root name server addresses. A validator could detect attempts to alter the IP addresses associated with root server identities and prevent such queries from going to a “spy-in-the-middle.”

However, signing the root-servers.net zone also introduces some potential problems. As demonstrated in RSSAC028, the size of a signed priming response can be significantly larger than an unsigned response, depending on the server software in use and the value of the DNSSEC_OK flag in the query. In that report, BIND version 9.10 returned signed priming responses exceeding 3800 bytes. In most networks this results in a UDP message broken into three fragments. It may be undesirable to rely on working fragmentation reassembly for priming queries.

RSSAC028 further showed that other software (Unbound, Knot) generate more reasonably-sized responses. RZERC may want to further study and recommend limits on priming response sizes.

Another aspect worth considering is how recursive name servers respond to bogus priming responses. In other words, do they actually check signatures and discard data from responses that cannot be validated? Initial investigations by Verisign indicate that some do, and some do not.

Lastly, it may also be worth studying and documenting the interactions of the .net zone between the root zone and the root-servers.net zone. Since root-servers.net is a second-level domain, validation requires the following RRsets:

- net DS record(s)
- net DNSKEY records
- root-servers.net DS record(s)
- root-servers.net DNSKEY records

This effectively changes priming from a simple, single query/response transaction into something more complex, requiring multiple queries to multiple zones.