

RSSAC Caucus Statement of Work for Studying Modern Resolver Behaviours

07 August, 2018

Background

The DNS is composed of authoritative servers, recursive resolvers and stub resolvers. Recursive resolvers are largely responsible for handling client traffic through a multi-step resolution process that involves sending queries to potentially many authoritative servers. There are numerous DNS RFCs that have been implemented in many open source and closed source products. However, community consensus regarding limited analysis performed on DNS packets received at authoritative servers is that the operational behaviour of these recursive resolvers is not well understood. It is additionally unclear if recursive resolvers treat the Root Server System (RSS) differently than other DNS zones in the Internet. Caching in recursive resolvers sometimes appears to be too short, while at other times too long.

This proposed work effort will seek to study and further understand the behaviour of existing deployed software and recursive resolvers through study of both code bases and available datasets.

Scope

The RSSAC Caucus Resolver Study Work Party will be asked to:

1. Provide recommendations about further work on the DNS protocol that is needed to better optimize deployment of DNS recursive resolvers and authoritative servers to improve performance and efficiency of the RSS.
2. Provide recommendations about potential changes to implementations and deployment configurations to better optimize deployment of DNS recursive resolvers and authoritative servers to improve performance and efficiency of the RSS.
3. If required, craft advice to other ICANN constituencies or the ICANN Board about any actions that may benefit the stability of the DNS and the RSS.
4. Document all analysis and simulation findings in a final report to be published as an RSSAC document

To undertake this work, it is expected that the work party will likely need to undertake these tasks:

1. Analyze DNS resolver network traffic and behaviour to better understand how they operate as they interact with authoritative servers generally and the RSS specifically in terms of preferred root server selection.

2. Analyze DNS resolver and authoritative server code bases and perform (ideally repeatable) simulations to further extract a model of how modern resolvers implement caching and priming of the root name servers.
3. Analyze DNS resolver code bases and perform (ideally repeatable) resolution simulations to further extract a model of how modern resolvers choose which authoritative server for a given zone to query.
4. Analyze DNS resolver systems using multiple resolution instances (with potentially individual or shared caching systems) to understand how they interact with the RSS. (e.g. google, cloudflare, quad9 type systems)

Deliverable

A *final draft* to be voted upon by the RSSAC before becoming an RSSAC publication.

Date of Delivery

The final draft must be submitted to the RSSAC no later than a year from work party formation. Submission prior to deadline is welcome.

Guidelines

The RSSAC requests that this work party first choose a leader, if the leader is an RSSAC member they will report progress on this work to RSSAC as appropriate. If the work party leader is not an RSSAC member, the work party's shepherd will report on the progress of this work. Fred Baker has agreed to be the shepherd for this work.

In the event that the deadline will not be realized, the work party should inform RSSAC immediately and provide details of the work that cannot be completed by the deadline. RSSAC support staff will assist the work party deliberation of the work, including setting up a mailing list for the work party, arranging and supporting regular teleconference calls, taking notes of meetings, and drafting background materials if needed.

References

<https://www.dns-oarc.net/oarc/data/ditl>