

RSSAC Caucus Resolver Study WP Meeting | IETF103
Tuesday 06 November 2018 | 12:45 - 13:45 UTC+7 | 05:45 - 06:45 UTC

ATTENDANCE

Fred Baker, Paul Hoffman, Kazunori Fujiwara, Wes Hardaker, Paul Muchene, Mohit Batra, Avri Doria, Harald Alvestrand.

Staff: Steve Sheng and Mario Aleman.

DECISIONS

- Fred is the shepherd from process perspective, Paul is shepherding the content.

ACTION ITEMS

- Mario to share a Doodle poll for the next teleconference, coordinate with Fred.

TASKS FROM RSSAC AND QUESTIONS TO BE ANSWERED

https://docs.google.com/document/d/1YRFhnJMDrXqpRnVtjZ8FkVUIyg_zEma_EqsuhM0awK8/edit

SUMMARY NOTES

Statement of Work

Wes: Statement of work. The real goal to figure out what the resolver in the plant actually do, instead of what we think they will do. How they deal with cachings? Etc. These are listed in the SOW.

Wes: These may lead to recommendations to resolver default behavior, the DNS protocol, the root server system, to ICANN.

Wes: The data: DITL data, code base of DNS resolvers, including simulations to understand how they interact with the root server system, and collective resolvers behaviors (e.g., query to 8.8.8.8 result in answers with random TTLs)

Fred: I would like to walk out today with more questions to be answered. Geoff has his own research and made some comments at ICANN. (see google doc)

Paul: ICANN gives a list of questions, how tightly should we do that?

Fred: My understanding, we are interested in anything that might would work. But more specifically the answers listed in SOW.

Wes: We don't want to limit the Caucus the work, but make sure to answer the original question. If there are expansion of the scope, we should let RSSAC know.

Question 1

Fred: Looking at the list, you had an issue.

Paul: Looking at public end resolvers, such as DITL data. When you look at DITL, you see the query, but you don't know if they are asking things all at once, or the requery. What are you looking? Because the DITL is two days and the TTL is two days.

Wes: We had two DITL for more than two days around KSK roll. We had prior caucus work that the TTL is two days, the query rate is much smaller. What's the standard bound on cache times?

Paul: For DITL, we had one root server that is not joining, also some of the data the IP address is anonymized. Also you don't know if the resolver is acting as a forwarder. I looked at L-root data and found this to be the case.

Fred: Dumb question. Could L is getting them from successive instances?

Paul: No, Roy is our data scientist. We are pretty sure of that. For KSK we are seeing this. So one thing may come out of question 1, can we seagate forwarding traffic versus organic traffic?

Wes: I think it is likely for all the section, we should come up with questions we would like to answer. It is highly unlikely we'd get to all of them due to data and other limitations. If some of the questions are so important, we may recommend longer period of DITL collection.

Paul: Maybe a longer collection for just root NS and DNSKEY data.

Wes: My colleague have code and data to look at aggressive resolvers. He is applying to Caucus.

Paul: So this is a broad stroke (or list) of questions.

Question 2:

Paul: The three ones I come up is what does it to for initial priming, when does it reprime.

Fred: Dumb question, we had a new RFC from Paul and Davey. He is asking we should abandon that practice because of fragmentation.

Paul: One of this will come out of this test bed is that we get surprising results. That we need to look into it further.

Paul: Infobox may tweak the result. Also Yeti is v6 only, that will have a big impact. One of the questions I want to understand if v6 interface is used at all.

[]: Redhat/Centos is running an antique version of BIND.

Paul: This would be part of it. So three questions: 1) From code bases and configurations, determine initial priming practices, 2) From code bases and configurations, determine caching of root NS, and 3) From code bases and configurations, determine caching of root DNSKEY.

Paul: My picture is that we will have a testbed. Set up a VM, and a script of different tests. That's my picture, and I believe we can do that. So we might need to prioritize question to finish this within a year.

Paul: If I do the testbed correctly. The results will be reproducible.

[]: You mentioned test bed. This would allow for what if analysis.

Paul: Yes, it would. Through changes in root zone TTL in the testbed, we can reduce the wait time to study the initial priming.

Wes: The reality is, looking at existing code is more important.

[]): We had resolvers doing ...

Fred: Let me comment on your question about additional root servers. It would be nice from simulations to see if it makes sense to reduce the number of root servers, or adding to them.

Paul: I like simulations, and in round two we may be asking RSSAC for more questions.

[]): Are you asking?

Paul: I don't know. I am thinking just a cache dump.

Paul Muchene: How soon the testbed will be available for simulations?

Paul: I am thinking to have the design to the mailing list before Christmas. I would be disappointed if it is not ready by February.

[]): My colleagues would like to get involved in this study, but he is not part of the RSSAC Caucus. What should we do?

Fred: We are stopping the work party list, and using RSSAC list instead. So you can search the archive.

Paul: Or your colleague could consider join the Caucus.

Wes: The caucus list is a readable archive.

[]): I think software developer knows this. So asking them may also help.

Paul: I agree, but we should also test it.

Question 3

Paul: I assume that if a resolver has been turned on. It primed. And it sent a query to L-root. I sent a second query for .net. Does it send to L-root? I call that sticky query.

Wes: We did a recent study at that exact question. They did determine it is 4 or 5. I will send the paper.

Wes: The root server operators don't necessary care about that data.

Paul: My question is whether the data would be helpful in providing operation advice for root server?

[: Maybe the question we should look at is the stickiness of the last query.

Paul: That's a good suggestion.

Paul: When does it refresh root NS set and for TLDs. How different is the root from TLDs?

Wes: I don't know which code basis does treat root and TLD differently.

Question 4

Paul: We could reproduce Geoff's stuff. How many resolvers ask queries over v6? How many of them do DNSSEC?

Paul: Geoff has been keeping consistent numbers. The third one is whether we can actually decide how many are doing the equivalent on aggressive NSEC?

Wrap up

Fred: I would encourage people to go to your google and modify the questions?

Wes: how detailed should the question should be?

Paul: As detail as possible.

Fred: We make sure we get the questions in that we think needs to be addressed. Let's finish that.

Paul: In the parallel, I will be working on the testbed.

Fred: When do you want to meet next? F2F?

Fred: How about in a month?

Paul: Before 15 of December.

ACTION: Mario to share a Doodle poll for the next teleconference, coordinate with Fred.

DECISION: Fred is the shepherd from process perspective, Paul is shepherding the content.

Paul: For those who are on the Board, you can read the archive or join the caucus.

Avri: Thanks. We have an active chair from BTC, so we are tracking things.

Paul: We may have a update to RSSAC before Kobe.

Notes from the Adobe Connect:

Discussion of Tasks from RSAC, first point: it's not clear how to do this without literally intercepting resolver traffic. How do we do this? Wes: derive from DITL data.

Can we separate forwarded traffic from "organic" traffic?