

RSSAC Caucus Resolver Study WP Teleconference Call
Thursday 06 December 2018 | 18:00 UTC

ATTENDANCE

Fred Baker, Japp Akkerhius, Paul Hoffman, Kazunori Fujiwara, Mohit Batra, Paul Muchene, Ken Renard, Wes Hardaker, Benno Overeinder, Shinta Sato, Joao Damas, David Lawrence.

Staff: Steve Sheng, Mario Aleman.

ACTION ITEMS

- Staff to send a Doodle to schedule the next teleconference
- Joao Damas to share a document related to the work and experiments with resolvers
- Paul Hoffman to provide an update of the different tests on the next teleconference call.

NOTES

Fred: Let's discuss with Paul and Geoff's lab.

Mohit: I made my comments on the list a few weeks back, let me know if you have any questions.

Paul: We are totally open to folks to WP and Caucus to suggest new things, ways to interpret. The research task is split into two categories: 1) what does current resolver software do in current environments (done in a lab) 2) what do we see in the wild, how can be better testing. Some are more testable than others.

Paul: Let me focus on the testbed.

(Paul showed a design of the testbed and explained the testbed.

Paul: I have set up (or planned to) three virtual machines. 1) Authoritative name servers that will do the testing. 2) resolvers that we are testing, we are testing all the popular resolver software. If people want to test a resolver not run in the VM, but as long as it can connect to the interface it should be ok. 3) the gateway. Slowing down responses from some of the servers. We will run our own root, push in a different trust anchors.

Fred: I would see this is how this is setting up alternate root.

Paul: Yes, alternate root is not difficult to set up. But we are in private network, so this would not leak. The server vm does not have direct connections with the internet.

Wes: It not terribly complex to set up alternate root. Yeti does that.

Wes: Overall, I see this architecture pretty good. For the gateway what are you using?

Paul: I am using ip table.

Paul: I am hoping TC would work, Wes could you send me pointers on other gateways?

Wes: yes, I can.

Wes: For the vm, are you running 13 vms?

Paul: A single vm running 26 addresses.

Wes: Ok, I can't think of any downsides of this set up.

Fred: What are you using to hide the IPv6 addresses?

Paul: I have not gotten that far yet. I might come to ask you Fred.

Paul: I am re-using the design our group have. So I am get this set up in January.

Paul: ok, has Geoff come on it or not?

Joao: What are you asking Geoff about?

Paul: Just his design of the test set.

Joao: I can speak to that. We use experiments that run javascript code throughout the world. What you can do with javascript in browser in browser is limited. We design names that are unique for each experiment. Each of these names when the client tries to resolve it has to go through a certain numbers of DNS hurdles. We have experiment where a name is only resolvable in A addresses, or only resolvable in AAAA. We have experiments about preferences. That's how we do the measurements. The interesting part of this is the DNS component, where you usually use the resolver that people are using. Some are using ISPs resolvers, public resolvers. Then we see which resolvers the clients is using, we can associate a client with a resolver.

Joao: Sometimes we design experiments specific to the resolver. By providing experiments with a valid DNS signature, or break the signature and see what happens.

Joao: [not captured here]

Joao: The DNS signal, the one involves the resolver, is very good quality. We can reach millions.

Paul: Do you have a write up of basically what you told us and we can share with the work party.

Joao: there is a presentation, and some texts. I will make a note and send to the caucus list.

Paul: Related to this, specifically one of the things we are trying to do is have the experiment repeatable by the researchers. Would you be able to add a little bit more detail on those?

Joao: Yes. The code base is on Github. The base idea (EVL – DNS) is a platform that is easy to modify to run these experiments. If you look for Ray's github, it is there.

Joao: Other than that is just parse the data, Geoff uses byson.

Paul: I think it would be helpful.

Joao: If you look at the EVL – DNS code, you give it zone files, and something like wildcard answering machine. The breaking of the DNS signatures is trivial.

Joao: Yes, we can write more details, but the code is in there. It is C++, not C.

Paul: Great, if we can get that from you folks. I think that would be helpful to folks who want to produce.

Paul: other folks have questions for Joao?

Paul: Hearing none, I think those are the two parts of the experiments. Over to you.

Fred: The other thing is Mohit sent a note. Do you have any thing that we have not touched on?

Mohit: Not no. I will go through the minutes and see.

Paul: I can respond to some of your comments.

Paul: your first questions is we should use a common recursive resolver and authoritative solver. Yes, that's what we are planning to do for recursive. I wasn't thinking about authoritative with different versions, but it is possible with the testbed. Does that answer your question?

Mohit: What we are doing google public DNS resolver? How do we simulate google public DNS?

Paul: No, we cannot. It is proprietary and it is constantly under change. If we send 100 queries to 8.8.8.8 from different locations, we would have different results.

Paul: We might look at the top 5 public resolvers. If you could help arrange that that would be great.

Benno: I remember Wes during IETF meeting, discussed an analytical model, and describe how resolvers behave.

Paul: assuming we have time, if we can fit in a one-year deadline of this project, what he discussed would be useful in general.

Paul: Google is hesitant to say, because people will be attached to what they say, where things are changing over time. If we can come up with a list of expected resolver behaviors, and test this against the testbed.

Mohit: If we can also do a list of researcher papers?

Paul: can you test it out?

Mohit: I can try to find out.

Fred: we've got another 13 minutes. Any other points to raise?

Fred: hearing none, I will give you the rest of the hour back.

Fred: when do you want to talk again?

Paul: I'd say we schedule a meeting in 5 weeks from now. We may have data.

Fred: we will send out a doodle poll to make sure people are available.