

RSSAC Caucus Resolver Study WP Teleconference Call
Monday, 25 March 2019 | 6:45 UTC

Attendance:

Alejandro Acosta
Fred Baker
João Damas
Kazunori Fujiwara
Paul Hoffman
Geoff Huston
Darren Kara
Mallory Knodel
Dave Lawrence
Paul Muchene(online)
Benno Overeinder
Anand Raje
Ken Renard
Shinta Sato
Petr Špaček
Willem Toorop

Staff:

Andrew McConachie

Actions:

Staff to send a Doodle for the next meeting sometime end of April or beginning of May.

Notes:

Paul Hoffman: Gives overview of Caucus Work Party schtuff.
Gives overview of RSSAC Caucus.
Gives overview of the work party.
Gives overview of testbed.
Gives more detailed info on testbed and how it works with virtual machines.

Willem Toorop: There is a testbed from CZ.NIC called Deckard and it's very good for testing these kinds of things.

Petr Špaček: We fake some system calls and use fake sockets.

Paul: The reason, that's not very transportable across control hosts.

Petr Špaček: That can't be measured across time.

Paul: We're not faking time. We have our own root and we crank down the TTLs to one minute. I thought it was a container issue. In this layout, as long as the thing you want to test. We could put hardware in there. But CZ'd Knot resolver works better from containers than from source. Having said that, I looked at Deckard a while ago and would like to look at it again. There shouldn't be a restriction on how to run the test. It has the latest BINDs and the idea is that the tooling should work. If we have to get a bit manual with that, we can swap in another container if necessary. The most important thing is to be able to control the gateway while also controlling the test.

Now I'll walk you through what currently works, so hopefully you jump in yourself and help out. I based this on VirtualBox since it is free. You get a copy of VirtualBox and start setting up the repo. Once you have the copy of the repo is you setup two VMWare images and get a copy of the ISO for the server box and for the recursives.

Dave Lawrence: Why choose FreeBSD?

Paul: With Ubuntu there are tools to add delay, but it can only be done by interface. Not by IP address. Linux has an optimization that unfortunately fixes this. FreeBSD works better for this.

Once you are done with setting this up you have the whole setup. If you go back to the repo and look at build-config.json, you will see the resolvers that we currently have. So we could add many more. The gateway is all setup and ready to go. The server side is also setup. At this point it's ready to go.

What we don't have is tests. And I would like to talk about that, but does anyone have any questions...

The first test I think would be best is the priming test. Given the fake root zone we're using has 60 second TTLs for the DNSKEY set and the RRSIGs. What should happen is that it should itself reprime. If it doesn't then we send another query for something that is not in its cache. Basically that's the test I'm thinking of starting with.

Did you guys try that with Unbound and Deckard?

Willem: For Unbound we use test-bound, which is what Deckard is based on.

Paul: Did you test priming?

Willem: I will take a look at that. I don't remember.

Paul: That would be interesting information.

I also want to acknowledge that Fred is sitting here. I'm not doing v6 yet, mainly for selfish reasons. What I would love is for someone who is v6 interested to sprinkle v6 everywhere and test with that. One question we have is how does a v6 resolver behave differently than a v4 resolver? happy-eyeballs?

Willem: What do you mean with happy eyeballs? For DNS it doesn't matter.

Paul Hoffman: Is that true? Is that how you looked at it?

Willem: v4 and v6 are treated exactly the same.

Paul Hoffman: OK. It would be interesting to know if older resolvers behaved the same way. I found out this weekend that knot expects people to always keep their resolver upgraded.

Any other questions?

Paul Muchene: I haven't had time to play with it that much. I was hoping to get time before ICANN in Marrakech to play around with the test code and give you my feedback.

Paul Hoffman: That would be great.

Willem: I just grepped for the word priming. Unbound has 5 test scripts for priming and 3 of them are also in Deckard.

Paul Hoffman: If you could send those around that would be great.

Willem: That's also configurable how proactive the priming is.

Paul Hoffman: It would be interesting to know what the defaults are, as well as for Knot.

As I said the 2nd type of testing is testing resolvers in the wild. Geoff and João are familiar with how they do their testing. One of the things that RSSAC asked for was for the testing to be reproducible. APNICs testing is not reproducible.

Geoff: It's not the Google ad network, it's that you need the ad to go to a 3rd party site. And that is a no-no. We got certification to do this with Google that would be difficult to get again. Lately we've been looking at what part of NO is difficult for the DNS. The issue is not resolver behavior. It's the way that many users sit behind collections of recursive resolvers. And your test bed doesn't do that. Because of the load distributors these are incredibly different. Some of them actually act like amplifiers. There is a 2nd order effect that we're seeing, that you can't see. Sometimes the behavior here is bizarre, unique and often unpredictable. This just adds to the soup of what goes to the root, cause analysis is tough.

Paul: RSSAC did not ask us to do cause analysis. We now now to not expect a simple answer to a simple question.

Geoff: Google get really annoyed about the domain names we're allowed to query these days. The original testing we did with generic TLDs is now not reproducible. Google is super finicky now.

Paul: For reproducibility you have blazed a trail that we might use, or not. We may do it differently in the future. What appears to be a single unique query, the same query happens 3-4 times.

João: At least.

Paul: It's hard for an RSO to tell that from their perspective. It may be the system you've setup is not reproducible. If people here are interested in doing this we may be

Geoff: To my mind reproducibility is less important. It's more interesting what the questions are and what makes sense. There is a unique root behavior, and other authoritative are different. The query traffic at the root is very different than other auths. The questions are interesting, and we are interested in knowing what is interesting to root servers. What questions would be interesting to the root servers.

Fred: We're not limiting this WP to this.

Paul: If we can do a test system that can answer lots of questions that would be great. Right now we don't have anyone to work on this.

João: What Geoff is saying is that we have this thing that is up and running. What questions are interesting for the RSO? We can runt tests.

Paul: One that was at the top of the list is the percentage of recursives that use DNSSEC validation. I would like to give the WP more time to setup its own parallel systems. Given that there is one setup that works, what does RSSAC want?

Who here is interested in pursuing resolvers in the wild?

Willem: I am spending time on this with RIPE Atlas. I've been running about 100,000 different tests since April 2017. With Roland van Rijswijk and Morritt Miller with specific queries.

Paul: Would you be interested in expanding that beyond RIPE Atlas?

Willem: RIPE Atlast is still good to get inside the diversity of the Internet. It gives a unique user perspective.

Paul: Is that of interest to you to give it to the WP?

Willem: Yes.

Paul: We are supposed to produce something by the fall. And take whatever Geoff or Willem has and put it into the document. Is that something that you are interested in working on?

Willem: Yes. We are also doing similar measurements with Illuminati. It's a tunnel broker. It's a commercial VPN broker. If you don't pay for the system then you are an exit point. If you have a paid subscription you can do web page fetches. and you can ask for a million different vantage points. We have a subscription and we use that as a probe and we have written that up. TJ from Rochester Institute of Technology is doing this testing.

Paul: Please send this all to the list. I'm assuming we would do that here. And maybe by doing this we find more interesting testing to do as well. Any other thoughts?

Kazunori Fujiwara: Capture may be interesting as well. Looking at DITL and trying to find interesting results. We are doing some investigation into DITL and should have the research published within the next few months.

Fred Baker: I believe most of the operators are doing a 72 hour collection around the removal of the key.

Paul: They didn't because they thought it wouldn't be interesting.

Dave Lawrence: But then it was interesting :)

Paul: If other folks are doing things with DITL data that is interesting. I'm going to be doing something about what does a recursive do with a single root server over the course of a few days. That might give us other questions that are interesting and it may spur interest in the RSSAC.

Fred Baker: Something that the RSSAC observes with these work parties is that they go off on rabbit trails. ANd that's fine, but also do the original thing as well.

Petr Špaček: In general we have a guy at CZ.NIC that is looking at past research and looking into what is out there.

Paul: That would be great to send in. Please do. It would be great to get other people interested in this research.

Paul: I believe that is it then.

Andrew: Staff will send out a Doodle poll for a meeting a month from now. End of April, beginning of May.