

RSSAC Caucus Resolver Study WP Teleconference Call
Wednesday 03 October 2018 | 14:00 UTC

ATTENDANCE

Fred Baker, Kazunori Fujiwara, Wes Hardaker, Paul Hoffman, Paul Muchene

Staff: Andrew McConachie, Steve Sheng, Mario Aleman.

ACTION ITEMS

- Wes to send some research papers to the list.
- Staff to schedule a meeting in Bangkok.
- Staff to investigate Google drive access for work party members.
- Staff to create a single document for output as well as brainstorming.
- Paul Hoffman is the leader.
- Fred Baker is the shepherd.

NOTES

Fred: My agenda this morning is to accomplish 2 things. Pick a leader first and discuss the scope. Who would like to be a WP leader?

Paul: I would be willing, but I am expecting to also contribute quite a bit on the test bed. Per my experience on other WPs one can do both.

Wes: I'm sort of in the same boat. I could go either way. However, if Paul can I think that would be fantastic.

Paul: It seems like the doc might follow the test bed. I think the scope will mostly follow the test-bed, setting up the test bed will be important. As we have a test bed that we can use, people might ask; why can't it do X? And this will allow additional tests as well as additional things in the doc. As long as I get nudges from folks to tell me to also include certain things in the test bed.

Wes: I suspect there will also be things in scope that will not fit in the test bed.

Paul: And it would be great if multiple people could contribute to the document as well. We've gotten better about Google docs.

Wes: That was one of my questions, do people like Google docs?

Fred: I favor Google docs.

Paul: I'm fine with doing this. We do have a smaller number of people on the call than expected. If someone not on the call wants to be leader that's also fine.

Fred: Kazinori do you have anything that you would like to add?

Kazinori: I agree that Paul could be leader.

Fred: Looks like we all agree that Paul could be leader.

Paul: OK. Let's assume I will be leader unless someone else jumps in the way.

Wes: Let me walk through the statement of work. There are really 2 aspects. There are resolvers that we know about, and resolvers that we don't know that much about. Some resolvers never actually cache. There are some resolvers that don't conform to how we expect resolvers to behave. Hopefully, we can look at DITL data to determine more about those.

Paul: Yes, I agree on both parts. One thing we have discovered is that you can look at how a resolver is distributed, but config greatly affects things. Some people setting up BIND as a validating recursive DNSSEC resolver are putting the public KSK in the DLV section. And some of that creates weird results. So some of this stuff will be speculative about misconfigurations.

Fred: So some of this will be speculative about what people might be using.

Paul: I don't know if we can do the fingerprinting. Some resolvers try to actively look like other resolvers. We can get code of all the various versions and build it into the test bed. I don't believe we'll be able to remotely say, "That person is running this version."

Wes: There was a good paper I read a while ago about detecting resolver pools. Some resolver pools have a shared cache. So it's hard to know where the queries come from. Not only to doc our findings but also to recommend if the protocol needs to change. Identify areas of code that would be good to read.

Is there advice that we should craft to other ICANN groups? Are there suggestions we could give to the Board?

And also, to document everything. We will likely want at least 1 Github repository.

Paul: We have an internal only test bed for stuff like this as well. For what we expect to happen after the rollover. So we've got that started. It is my intention for us to not only doc the test bed but to make it easily distributable. Because we got a little beaten up at the last Board meeting about lack of reproducibility on studies that we are basing decisions on. Not just to have a test bed, but so someone can reproduce our results. Us having strangers able to reproduce our results will enhance our credibility.

Fred: Probably. And it seems like it would be good to reproduce in multiple locations.

Wes: I know Roy did some work in this area as well.

Paul: And I did too as well. A while ago I created a test harness that was only ever used by Duane Wessels. Part of this is to make it, then make it simpler. Gee, that result looks funny and I will try and go do this as well. Some of RSSAC members may want to hand this off to members in their operational teams. The earlier we can make this available would be better.

Wes: Any idea of how soon you can push your test bed into something public?

Paul: End of January hopefully. A bit busy now. Maybe December. This is not me working on it now, the design is sound. Instead of having 13 virtual machines we have it all in one giant VM with 12 instances running on separate IP addresses. So then it gets easier for someone to try it on their own laptop. I would really hope we could have this in January.

Fred: So you want to put multiple resolvers in the same VM?

Paul: But not running at the same time. You start the resolver, you do stuff. You shut it down. You start up the next one, then see what it does. One resolver at a time with lots of tcpdump and lots of logging. It could be one resolver with 13 different config files.

Fred: And we can blow it away each time.

Paul: Some of the popular resolvers can't be built by hand that easily. If you want different versions of them. Building them from repos mean you only get the latest versions. And we also want Windows.

Wes: The sooner we can do everything in public the better.

Paul: We're still on page 1 and I have questions about page 2. We've been talking about the test bed for known resolver software. You also mention "resolvers in the wild" on bullet #4 second list.

Wes: This SoW was a collaborative effort. We have a fair amount of leeway on how we want to approach this. One of our first goals should be to define a concrete list of what we want to answer. Which of those can we answer? The importance of the bullets is really up to the work party.

Fred: Especially the idea of studying resolvers in the wild.

Paul: Yes and No. Geoff Huston does study resolvers in the wild. We're using his results for the KSK rollover. He has a full test bed for doing that. And there hasn't been a public description. It is parallel to everything we've been talking about. They do look at stuff like amount of V6 or DNSSEC. They're listed here in #4 but not heavily discussed. I think it's fine to do parallel tracks, but we need to know that.

Wes: Hopefully Geoff can contribute. There is a lot of intellectual property that makes it unreproducible.

Paul: And he also doesn't control where the samples are taken. For example, some weeks all the queries come from one country.

Wes: It was hard to not say, we want X, Y,Z. When you're diving into a study like this you don't want to over prescribe the result. It's hard to do when talking about the root servers, and one thing we don't know is how resolvers treat the root differently than other auth servers. And there is to some degree because of bootstrapping. Multiple resolvers with shared caching is one of the biggest unknowns. We shouldn't reproduce work that has already been done. Some of this stuff is really complex. To answer Paul's question, this list of 4 questions lets not limit ourselves to only these 4 things. The WP has a lot of leeway to determine what is important.

Fred: This all sounds pretty familiar.

Paul: As we move forward then we can see a bit more about how we can do that. I have been interested in trying to get reproductions of his work.

Wes: We really want to study the resolver behavior throughout the root system as much as we can. Both what we see in the world and code bases. John Heideman had an interesting study recently with Anycast where some nodes were constantly flipping back and forth. That clearly shows that even understanding how resolvers pick their destinations routing comes into the question.

Fred: Is that picking different instance of the same RSO, or different identities?

Wes: Different instances of the same RSO. Which makes it harder to study as well. I will send this paper as well.

ACTION: Staff to investigate setting up a Google Drive for the work party

Fred: I have a comment on the SoW and creating one paper on what we produce. We might want to create two papers.

Wes: I don't think there is any constraint on publishing just one paper.

Fred: Do the academics want something that could be publishable somewhere else?

Wes: I don't think we'll write that kind of paper. That's outside the scope of this particular effort. But maybe we could use the research.

Fred: When do we meet next?

Paul: My feeling is a month from now is probably the earliest. I think a month is probably sufficient. So at IETF in Bangkok.

Wes: I will be in Bangkok.

Paul: I expect Geoff will be in Bangkok.

Paul Muchene will not be in Bangkok.

Fred: Let's leave that Sunday afternoon meeting at ICANN 63 for the moment. But I expect the next meeting for us will be in Bangkok.

ACTION: Staff to schedule a meeting in Bangkok.

Wes: We need a space to do some brainstorming.

ACTION: Staff to create a single document for output as well as brainstorming.

Adobe Connect recording: <https://participate.icann.org/p9fxrnite57/>