

---

MARIO ALEMAN: Welcome everyone to the RSSAC Caucus Resolver Study Work Party Teleconference Call on Monday, the 14<sup>th</sup> January, 2019 at 1600 UTC. Many of our guests are new to this teleconference, so please welcome and I'm going to turn the microphone back over to Paul to start the call.

PAUL HOFFEMAN: Actually, I'm going to turn it over to Fred. Fred, why don't you give us a little bit -- since there are some new people, give people a little bit of a background before I start diving into the work that is happening.

FRED HOFFMAN: Okay. You should have received the -- on the Caucus list, should have seen the link to the Statement of Work. The project is in essence -- pulling up the ICANN WIKI Page which is in that email, the question here is kind of, how do resolvers work? How well do resolvers work? The RSSAC of course has its primary mission relating to the root itself but an important part of that delivery system is then to go to the Google's and the open DNS's and the whatever of the world and then actually deliver that information and the TLD information to the world at large. We want to know how they work? We want to know how well they work. We want to know what configurations would be appropriate to build those?

We've had three meetings so far; one was last October, one happened at the IETF Meeting in Bannock and so for those that were there, it was face to face. Some were online and some were elsewhere, and then we

---

*Note: The following is the output resulting from transcribing an audio file into a word/text document. Although the transcription is largely accurate, in some cases may be incomplete or inaccurate due to inaudible passages and grammatical corrections. It is posted as an aid to the original audio file, but should not be treated as an authoritative record.*

---

had a meeting in December, so this is the fourth meeting. The project plan so far has been to try to follow basic scientific practice and document a test bed that would allow us to put that together. Paul has been putting together a test bed at ICANN and has put some of it, the start of it, into GitHub, I'll have him discuss what he's doing. Jeff Huston has been running a bunch of experiments and posted a note to the Work Party Mailing List last night, describing in general terms what he's being doing APNIC. I don't believe that was shared with the general caucus so some of us may have gotten it and some of us didn't.

At any rate, what we'd like to do today is have Paul discuss the work that he's been doing and Joel says he'll be on 15 minutes late. Jeff indicates that the time is pretty bad for Ken. Ken for Australia, he'll try to get on but he might not make it, we'll see. Hopefully we'll get some comments from that group as well. Paul, let me turn it back to you. What do you have going on?

PAUL HOFFMAN:

Just to be clear because Fred said "test bed" I want to be clear, we're going to have two sort of test beds. One is for testing resolver configurations in a lab, meaning we will set up resolvers and we will set up a lab environment and we will make them do things that would interact with the root and we will watch them. That one I'm calling the Internal Test Bed because there will be no real users there.

The second test bed that might or might not do, although we can certainly, hopefully make progress towards it, is what folks, actually Europe I forget who are calling it Resolvers in the Wild, that is tests that

we can get real users to run in possibly various ways, that would go through their own recursive resolver and then we can watch results from that. So, two very, very different things.

What I have this morning, let me paste the URL to the chat here for everyone, is let's first talk about the first of the two test beds, the Internal Test Bed. I sent out this link last week, haven't seen any traffic about it. Basically, I am setting up something that is reproduceable, hopefully anyone would be to do this on their own laptop or computer, would be able to do the internal testing. Where I stand now with it is, I've got the network outline and some of the bits going. It is certainly not useable at the moment.

Where we are at is, there are full instructions for how to us and we're using Virtual Box just because it's free and people seem to like it okay. How to set up Virtual Box for the eventual test bed and to start pulling down some of the resolver software. What is still left to do and I certainly don't think this will happen instantly but the major things that are left to do is, I don't yet configure the authoritative server that we are using for the root zone and again, we're going to have to build our root instead of using the real root for two reasons.

One is, we want to be able to be sure we're watching the traffic completely but more importantly is we're going to want to change the TTL's in the root zone so that one of the tests that RSSAC asks to do is, is not only look at priming but also at re-priming, so after a resolver is connected to A root. When the two-day TTL expires, how do they pick or do they pick another one? Do they start testing before then? For that, we don't want these tests to run for two days. We have to make

---

---

the TTL something more reasonable and then follow through. We're going to have to set up essentially our own root. That configuration needs to happen.

Certainly, one of the things that RSSAC wants us to do is to look at root selection based on delay times for an individual, so I have to put that into the router that I have in that diagram and that's fine, I don't expect that to be hard. In order to use our own internal roots, we have to have all of the resolvers have a modified root in its file, I have to actually come up with configurations for them. Again, I don't think that that will be that hard either.

Beyond that, we actually have to write tests. We did a set up like this in ICANN last year in testing for the KSK Rollover, so we have one design of tests where we have a certain cadence of queries going to the resolver in order to watch but, I'm hoping that this group will also contribute ideas for that. What is the best way to run a test that will get us the most information on the resolver's capabilities? I still picture that we can have this happening pretty well and maybe be into the thick of testing, depending on contributions from this working group, from the Work Party in two months, which is also when we are meeting -- actually I'm sorry, I've totally forgotten, are we meeting at the ICANN in Kobe or at the IETF in Prague, Mario?

MARIO ALERMAN:

Hello Paul, the meeting will be at the IETF most likely.

---

PAUL HOFFMAN:

Thank you. I'm sorry, I'm just looking at my calendar and they both start with the letter I and they're both in March. By the IETF meeting in March, I believe we, depending on contributions from the Working Group and again, I'm sort of assuming I'm going to be doing most of the work but I am -- from that GitHub Repo, absolutely looking for issues and pole requests and someone wants to contribute code or even better, fix code which is a little bit hand fisted, that's just fine, I would love to see this be a group effort. I think we will have a much better idea of how we are doing in March and possibly be already working on those tests for the Internal Test Bed.

For the Resolvers in the Wild Test Bed, it's a very different situation in that we are currently a lot of people are using the information coming from APNIC and as Fred said, we've just gotten the first of the more detailed description of what's going on from them. We might spend the next few months then looking at that description and asking questions, maybe filling it out more, having it filled more and then we as a group can decide, do we want to for the answers the questions RSSAC is asking, do we want to rely on the APNIC methodology? Do we want to create our own methodology that might use for example different data sources and such? Do we want to do a combination? That's all very open.

On the last call, some people talked about the idea of us setting up essentially our own methodology and running some tests if we can get what we consider to be useful data sources, that is to have a way to have the queries going from a large number of end users so that they go through a large number of resolvers. There's already been some research on this being done.

---

APNIC uses the Google Advertising Network and we could use a little bit more description of that for what their limitations? Can we do it better in some ways and such like that? There are other ways to get queries out to users behind resolvers. People mentioned in the last call that people are doing this with RIPE ATLAS, which has different properties than Google Ad Network. There may be other advertising networks that we can expand into. There may be ways of putting things on widely hit websites. Possibly in widely hit Java Script Libraries.

The black box of how do we get queries going is one thing that will certainly affect our results but then if have a methodology that we like, we can let other people contribute how they want. That work is just getting started, like as of last night that we have some to work on. We don't have to set it up as a single test bed either. I think there is definitely a desire for some parties to run their own tests, using their own data, that they don't want other people to see but they're willing to show their results. That's where we're at now. Why I don't I take questions on the test bed if there are any first, on the Internal Test Bed. Did anyone have a chance to look at the GitHub Repo?

PAUL MUCHENE: This is Paul here. I've actually taken a look at the Repo and I'm actually doing a set up right now.

PAUL HOFFMAN: Very good, thank you. And Abdalmonem, you have your hand up?

---

ABDALMONEM GALILA: Yes, this is Abdalmonem Galila. I'm sorry for being late. Just I want to ask; is this resolver test bed, will it be -- I use it for testing the ICB's resolver or my own organization's resolver?

PAUL HOFFMAN: Only for the open source resolvers or anything within the test bed. It would not be for any live testing. So this is -- I'm sorry, go ahead.

ABDALMONEM GALILA: You mean I could use this test bed environment to test the public DNS for Google public resolver, which they are using recently DNS for DNS queries? Would this work for that or not?

PAUL HOFFMAN: No, this would definitely not work for that. The Internal Test Bed is only for testing recursive resolver software that is running in the test bed itself, so things like [inaudible], Unbound and such like that. It does not send any queries out to the net at all.

ABDALMONEM GALILA: Okay, thank you.

PAUL HOFFMAN: The second test bed, the Resolvers in the Wild test bed very much would be used for that and again, we haven't gotten on that. For in the Internal Test Bed, it would only be for and again, the purpose for this really is to do tests that wouldn't really make much sense for Resolvers

in the Wild such as, there isn't a good way for us to test root server selection for even the resolver in your own organization, just because you'd have to watch for two days or whatever. The Internal Test Bed would be good for controlling all the variables, being able to crank down the time close enough and then get results that way.

If no one else has hit the test bed yet at the Repo, I think what I'll ask is that people who intend to test, would in fact pull down the whole test bed now. The test bed should build and if doesn't or the requirements are wrong and then I will send messages to the mailing list as I get through incremental steps. It would be lovely if I felt like I wasn't the only one doing this and so within a couple of weeks that there'd be people who would say, "Yes, it builds, no, it doesn't." You can send messages to the mailing list or open up issues in GitHub. If you have an issue with GitHub let me know, that would be great. Benno, you have your hand up?

BENNO OVEREINDER:

Question, I did look at your test bed, also just this hour. The scenario you just mentioned about the resolver's prime and re-prime, it might be interesting to describe, and I was also thinking, just today I had a brief talk with Ralf, one of our software engineers, about how we are doing the priming; it makes sense at least of the people here involved, the software developers, to describe in words how we think it works, how we think we have implemented it, and if you see it's reproduceable in your test bed, it doesn't make sense. Do we need to validate our own assumptions or validate the test setup?

---



PAUL HOFFMAN:

That's a very good question, thank you, Benno. It can be done in either direction. For the folks who have not looked at the test bed yet, we are not only testing the most recent version of any of these resolvers because we know that old versions are still running out on the internet, so we look at the last sub version of each of the major versions. There's five or six version of Unbound going back four or five years, the same is true for Find and such.

We could wait for you to describe what you think you're doing for resolution or we could run the test and you could try to -- since I'm not going to look in your code, you could try to then say, yes we believe that this is showing that we were doing what we thought we were or it would quite frankly be interesting to hear from software developers who say, we thought in version X, that we were priming this way but the test results show that we're not, so let's do some more intensive testing, specifically to see why there's a discrepancy. I wouldn't be at all surprised if the results of this test bed ended up surprising some software developers, particularly because some of this is based on configuration.

One of the things in the test bed that we will absolutely be able to do is, is for any given version of a resolver, run multiple configurations. As I get further along in developing the test bed, let's see but it would not hurt at all if any of the software developers on the call or on the mailing list wanted to start looking in their code and describing what they think they're priming and re-priming strategy is and therefore what they would expect results to be. That would be great and that will also help

---

---

inform us of what kind of test do we need to do. I would be very interested in hearing from Bind and Unbound and Power DNS, which of the user visible configuration options they think would possibly change the priming because then we can certainly test all of those options.

**BENNO OVEREINDER:** Thank you. We will come up with a description this week, next week. Andre was not able to join us, he was struggling with...

**PAUL HOFFMAN:** If the Bind folks can do that as well, that would be lovely and we're not in a rush but it would be good like I say, because as we start developing the actual test methodology, it might be for example one of the software packages says, we treat the root zone differently and we probe for the fast root this quickly as compared to any others or there might be a setting that says, if in our testing we find three that are identical we're going to round robin between them, then we'll know to add a test for a fourth, things like that. The more description we get of the software that's in our test bed, the better.

The other thing to be clear on is that even though we currently only in the test bed testing, the open source that we can pull down, I designed the test bed as best I can, so that if somebody has a piece of non-open sourced software, they can put on, even if it's hardware, that they can give addresses to and give the root file, it should work just fine in the test bed. I would love to hear from folks who have proprietary resolvers or specialized resolvers that we can't distribute what their test results are as well.

---

Are there any other comments or questions on the Internal Test Bed first?

UNKNOWN SPEAKER: Just with my [inaudible] team add-on, we are just finishing the next stable release, so I'm not sure the team will be able to spend that much time in January, but after that, we'll look into that. This is a very busy month for us.

PAUL HOFFMAN: Okay. Well, we like stable releases, don't let us get in the way. Moving on to the Resolvers in the Wild Test Bed, Joao is on the call now. Joao, do you want to talk a little bit about -- and you didn't hear, but Jeff only sent the message last night and only sent it to the internal mailing list, so do you want to talk a little bit about what Jeff sent and how we might work on this as a living document?

JOAO DAMAS: Okay. Can you hear me?

PAUL HOFFMAN: Yes.

JOAO DAMAS: Okay, good. First thing, could I request to be added to the Work Party so I don't get to the off site for things like this?

---

PAUL HOFFMAN: Sure, although aren't you on the RSSAC Caucus?

JOAO DAMAS: Yes.

PAUL HOFFMAN: Fred and Mario, we need to talk about maybe shutting down the Work Party mailing list so that doesn't happen again. We'll take that offline but that's an action item because we need to figure out soon the difference between the Work Party mailing list and the general mailing list.

JOAO DAMAS: Solution that works for everyone.

PAUL HOFFMAN: I think that was just Mario saying that he's marking down the action item.

JOAO DAMAS: What I would like to know is more of with regards to what Jeff said, is that along the lines of what you're looking for, the group in general? Do you want us to drill down more in a specific area?

---

MARIO ALEMAN: I'm sorry that is the operator.

PAUL HOFFMAN: I'm not sure who just left the call. Go ahead Joao, sorry about that.

JOAO DAMAS: Anyway, I was just asking, based on what Jeff sent and what we had discussed in the previous meeting, is along the lines of the detail that the party will be looking for or not. And then anything in the document can surely be a public thing. I mean, we'll describe what we do in any level of detail that's necessary and relevant.

PAUL HOFFMAN: Since only a few people have seen this, why don't I take it upon myself to -- I'll make sure that the document get's out to the full mailing list with those questions because, you know us, we can always ask more questions and if you folks are willing to have this be a document that get's expanded a bit, then as people on the mailing list ask for or as you say drill down on some of the specifics, I certainly know that for us, it would be very useful for example, to know on the validation test, exactly how your queries got matched up in both the authoritative server log as well as in your web server log. I think that there will be some drilling down that people will have.

Why don't we take this to the list first? Do you know if this document lives on a GitHub Repo or something where people can comment easily line by line or is this just something that you folks have in a text bed?

JOAO DAMAS: Right now, it's just a text file.

PAUL HOFFMAN: Okay, why don't we start with that? Why don't we start with the people can ask questions, once we get it out to the list, people can ask questions and then we'll figure out how we want to move it forward?

JOAO DAMAS: Sounds good. Happy to post updates of the document in the meantime as we progress, as we get the testing.

PAUL HOFFMAN: Great. Let me as a question that I think other people probably want to know too. Is this a document you put together just for us or is this something that you were intending to publish in some other venue anyways?

JOAO DAMAS: No, we put it together for this group based on things we had from previous presentations and so on. From time to time we do these recaps of why and how we do the experiments but not too often because people are interested but not that much.

---

PAUL HOFFMAN: Yes, I can understand that, having discovered that almost no one had looked at my Repo this morning either. We'll try to come up with the right level on that but I will make sure. Give us a few days to figure out the mailing list and such but then I'll make sure that it goes out to the full list, with a request for comments and further questions and we can see where to go from there.

JOAO DAMAS: Sounds good.

PAUL HOFFMAN: I think that's it from me. Does anyone have any questions either for me or for Joao for the APNIC Testing?

FRED BAKER: Joao, now I'm looking at Jeff's comment in email sent to the Work Party list last night, he said that he wasn't entirely sure what Paul was looking for or what the Work Party is looking for in terms of paper and so on and he had or you and he had put together this paper. Question, do you, having listened to Paul this morning, do you have a better idea of what the Work Party is looking for in a description of Resolver's in the Wild testing?

JOAO DAMAS: A little bit actually. Last time around we actually spoke a bit about the infrastructure and how we use to deliver the test to the clients and what we get in return. What I heard Paul say today was from the data

---

that we get back, how do we extract meaning? We can elaborate on that part.

FRED BAKER: Okay.

JOAO DAMAS: We get back scattering information from the clients, we can put together this means that and that's we did the experiment this way and that's why we test these applications.

FRED BAKER: Okay. Are you a position to discuss any of that today?

JOAO DAMAS: I can, yes.

FRED DAMAS: Okay, well, the other half of the agenda was to let you and Jeff do so. Let me turn the floor over to you and I realize I've got you a little bit flat footed.

PAUL HOFFMAN: Let me prime that discussion a little bit with something that you and Jeff and I had discussed a few years ago, is the way that you get results is you look in your authoritative logs but, for some of the tests you also, because they're getting back an address for a web resources, you also



---

look in your web logs. Can you just describe a day in a life of one of those tests, with what you get back and how you look in the logs?

JOAO DAMAS:

Sure. As you were just saying, we have two main components, one the client is running the test for us, let's forget how gets delivered, it just happens to get delivered and at the client. One is indeed a web server where the client is instructed to search a one by one side PNG, very small file that fits in there, that packet and then some. That is the olden way test of whether the client got there or not. This is also because inside the browser you are limited in what you can request the java script to do, manly for security reasons, which are good security reasons.

We want to test things you can do, ask the client to fetch a URL from somewhere else. That's what we do and if the client is successful in going through all the DNS obstacles, we throw in front of it, then it will get to fetch from our web server, we have a signature in our web server log. The web server signatures do have a potential problem, which is if the client aborts the test before has a change to fetch and there is always certain delay right, then we might not see that signature.

We have a bit of a control of a time limit so that after a few seconds, usually 10 seconds, if we don't get the result we will try too -- if we don't have on the client side the record of the different experiments having completed, the client will send through the server a different type of fetch, a different name will be fetched that includes query string indicating which experiments were able to complete and which were

---

not. We mostly always get a final sort of wrap up fetch and that allows us to look at the information.

Sometimes we don't because the web ad is aborted and then it never has a chance to complete. Most of the time that is still okay because when we are measuring the DNS we are really concerned about how the DNS works and the beauty of this is, once the client initiates the fetch, even if it doesn't complete on the clients web browser or mobile app, the query will already be initiated and the next query will have been issued towards the resolver and the resolver keeps working whether the client will still be there at the end to receive the result. It becomes synchronized and independent of the behavior of the end user.

Typical experiment has three steps, one is addlette gets loaded and the first thing it does it comes to our web server and fetches a list of experiments to run. Then the DNS process takes place, which is triggered by the client-side initiated fetches. Then the summary at the end is also sent. We know which client asks for the information in the first place in terms of IP addresses because the first thing we ask is for the configuration. Each configuration, each experiment has a set of unique string made by putting together the time stamp and some random values and sometimes an encoded address if it's puritan.

When we see the DNS coming towards, that same string, we know that this is a resolver doing work on behalf of these other clients that is running the ad. Then we see whether the DNS result succeed, if the browser is going to the end with it and we see the summary. In the validation for instance, we need to see all the aspects. We see first inside the client for further analogies down the log stream. Then we see

---

the DNS resolution, so for instance to test whether someone is validating or not, we ask for a URL that is unsigned which should always be fetched. One that is signed, validly signed, which should be fetched most of the time, independently of whether you are validating or not, and then one where we break the signature on purpose. These will be the people who are actually validating or just asking for DNSSEC information but not really doing anything with it.

And so, in the [inaudible] validly signed experiments or names, we get to see whether the client is actually doing something with them. If we see the web fetch right after the DNS resolution succeeds. We expect to see for instance in a typical case validation resolver, we will see the fetch of the unsigned, fetch of the DNS signed, but not the fetch of the invalid signed. If we see the three, we know we have a resolver that is asking for DNS site information but not really doing anything with it. If broken configuration, perhaps we only see the first one, when old keys are present that might be case these days because it fails a validly signed name.

Those are three steps. We know who asks in terms of who is the client, that also allows to map the client into an autonomous system and sort of give fair economy geolocalization. We don't go any further, we are not really interested in whether this person is in this city or this other city, we look more at the network level. Even address belongs [inaudible] prefix which is always [inaudible] number and that gives us a good indication of where the client is. Then we see the resolvers that are used by the client, the final we see the client coming or not coming back to us for the web page.

---

The different things that the client and the resolver manage to do or not, gives us new information about whether they are validating or not, whether they have correct keys set up, whether they are just fetching but not doing anything, perhaps passing it along, that could be also. There are things that we normally don't see and this is for instance the case where the resolver is not just a resolver but a set of forwarders that points to someone else. That's really hard to [inaudible] to the DNS, not mine.

We have a few [inaudible] that we're thinking about implementing in the future but normally cascading layers of forwarders are hard to see, you just see the other surface of a resolver farm or an ISP who decides that the resolver is actually going to just be a forwarded and points to Google for instance. Did I give you a decent idea? I can write it down so it's a little bit more clear.

PAUL HOFFMAN:

Yeah, so writing it down -- I won't speak for everyone, for me and I've heard you give this talk before, yes, hearing that is very good, having it written down would be excellent. Maybe it can be part of this document, the one that we'll be sending. I suspect that we will have a lot of questions on that because it would be nice to be able as you said at the beginning, it doesn't matter where you got the initial queries but it does matter what the java script code is. Even having some technical level like that would be useful.

Why don't you let us start with what you have and I think we will end up asking questions on almost everything that you said, does that sound

---

like a way to go? If you don't mine iterating over this for the next month or so? Or face to face meeting is in two months, why don't make that as a goal of, try to get it as clean as possible for that meeting?

JOAO DAMAS: That sounds good.

PAUL HOFFMAN: Well in advance, meaning people are going to ask you questions and I would love to Rev's on the list.

JOAO DAMAS: Sure. When is the meeting schedule, that's another detail I may have missed?

PAUL HOFFMAN: Mario, do you want to tell us? Tell you what, let's hold that question for a moment. Do other folks have questions for Joao on what he just spoke about? I'm not seeing any hands in the Adobe Connect. Okay, Mario, why don't you talk to us about future meetings? The face to face meeting is in two months, we might have another meeting in between such. Can you sort of walk us through that?

ANDREW MCCONACHIE: The Work Party will probably have another call next call, of course that's up to the Work Party but I'm kind of assuming given the cadence we've been keeping they're be another call next month. Then there will be a

---

meeting at IETF 104, of course there will be the full RSSAC Caucus meeting on the Sunday and the meeting for this Work Party hasn't yet been scheduled but it's typically early in the week at IETF and we'll send out a Doodle at some point, once we have a better idea of availability of rooms and remote connectivity and that kind of stuff and get that meeting scheduled. There will be a meeting of this Work Party, I'm assuming next month, teleconference and then there'll be another at IETF 104 in person.

FRED BAKER:

Andrew, do you know what day that's likely to be?

ANDREW MCCONACHIE:

I don't, it's really up to the Work Party. We haven't had enough Work Party meetings at IETF for me to have a big enough sample size to know what date would be best. In general, my experience has been that people have more time at the beginning of the week than at the end of the week, so maybe Monday, Tuesday or Wednesday but right now I couldn't really say.

PAUL HOFFMAN:

We also have to be sure that whatever we do doesn't conflict any IETF meetings that people want to go to, in specific, DN OP, Depriv, things like that. Andrew, do we have the capability of having these meetings in the lunch slot?

---

ANDREW MCCONACHIE: We usually do, yes. Verisign has been really cool about letting us use the room that they usually have. Hoping that they let us do that again 104, then we have a lot of flexibility and then there's always good remote connectivity there as well. I can do pop up remote connectivity so that's not as big of an issue, its more about finding a room and finding a lunch slot that isn't already taken with some IETF or ISOC activity. We'll send out a Doodle as the date approaches and the Work Party will hopefully arrive at a time that works for the most people.

PAUL HOFFMAN: The IETF has been better lately about having multiple open session rooms available and so as we get closer, we might want to do the Doodle pole early, so that we can get one of those, although they often block them out during lunch. Let's see what we can do and also talking with the Verisign folks and such like that.

I personally believe we have plenty to talk about in a month. That is that I will have made a reasonable more amount of progress on the Internal Test Bed and it sounds like we will have at least a bunch of questions for the APNIC folks on their designs that would be useful to talk about. I think it would be good for us to have another call in about a month.

FRED BAKER: Okay, about a month would be about the 14<sup>th</sup> of February. Let me just ask the collected crowd, is there a date that's going to be better than another in that timeframe?

---

JOAO DAMAS: On the 18<sup>th</sup>, so starting the weekend of 16<sup>th</sup> and 17<sup>th</sup>, both Jeff and I will be on the West Coast of the US, if that helps. [Inaudible] participate in the meeting.

FRED BAKER: Okay, so what are the dates? Is that the 16<sup>th</sup> through the 18<sup>th</sup>?

JOAO DAMAS: The Monday of the [inaudible] starts Saturday and Sunday the 17<sup>th</sup> I believe.

FRED BAKER: Okay. Yeah, the downside of that particular Monday is that it's a public holiday in the US. We can probably set it up for the 18<sup>th</sup>.

JOAO DAMAS: We could do it Tuesday as well?

FRED BAKER: That was kind of where I was going with the question, would Tuesday work? That would interact with [inaudible], right?

JOAO DAMAS: Yes, but finding time while at [inaudible] is kind of easier, at least it's not 3am for Jeff, for instance.



---

FRED BAKER: Getting Jeff on the call is an objective.

WES HARDAKER: -- start Tuesday at 9 Pacific, so actually you could do 8am Pacific on Tuesday.

PAUL HOFFMAN: Wes, can you repeat that?

WES HARDAKER: I can try. [Inaudible] I think normally doesn't start their morning sessions until 9 or 9:30 or something, so in theory we could do the same time, 8 o'clock Pacific on that Tuesday, that should impact.

PAUL HOFFMAN: Great, so Mario or Andrew, I'm not sure which one of you are now doing the Doodle poll but why don't we -- you might even do a one question Doodle poll of whether 8am Pacific time on February 19<sup>th</sup> works for folks for the next call?

MARIO ALEMAN: Sure, I can take that poll.

PAUL HOFFMAN: Great, thanks. Fred, do want to close us out? Ask questions?

---

FRED BAKER:

Well, it sounds like we've got some actions. We need to figure out where to put these text files that are being developed and Paul you have suggested GitHub. I'd like everybody in the Work Party to go to the same place for all the information and not a different place. Even if it was just a list of links to various things, it would be nice to go voice for the list of links. I've been talking with Mario in chat about the possibility of using the ICANN WIKI for that, we'll figure something out. I'll be looking for a Doodle pole on when to have the next meeting. We'll probably also end up doing something like that about the meeting at the IETF.

One day that I'd really like to avoid at the IETF meeting is Sunday. It seems like the world at large jumps on and okay, there's nothing on the IETF on Sunday, so its Sunday. Literally, I'm in meetings all day on Sunday and that's usually when the RSSAC and the RSSAC Caucus have their meetings at IETF meetings. We'll have that conversation in email and on the caucus list.

As far as the difference between the caucus list and the work party list, let me explain a little bit. At ICANN 63, the RSSAC had a bit of a conversation about the difficulty of getting people to actually do things, to participate in the Work Parties. I think at least part of that is procedural, they don't know how. What we thought about was instead of using the Work Party list, which we create, use the general caucus list and if the noise level isn't too high, that probably works.

---

The question would be, would there be people in the caucus that don't really want to hear about the work party? At least for the moment we're trying out having things on the general caucus list but I don't think that message is uniformly understood. Mario, Andrew and I need to talk about that. I think we've got an action for a call roughly the 19<sup>th</sup>, 18<sup>th</sup> or 19<sup>th</sup> in February and the plan for a meeting done at the IETF in March and we'll proceed from there.

One thing that would be really good, we talked today about write ups of the priming procedure, if people could just email something to caucus list, saying in the subject line use the word resolvers so that people can filter for that but then document whatever their code does and I'll be talking with Mario and Andrew and Paul about how to put that all in the same place, somewhere.

PAUL HOFFMAN:

Since that's related to the resolver test bed, I can create a text file there for that.

FRED BAKER:

Works for me. The point is, I want it to be accessible. Joao, we're going to need whatever you guys can develop to describe your test bed, your methodology and explanation on the note that Jeff sent out last night for the same reason that's Paul is doing it. Did I leave anything out? I think that's pretty much all the items coming out of this. As people put together their statements about what their resolvers priming technology is just post it to the caucus and we'll pull that together somewhere. Did I miss anything then?

PAUL HOFFMAN: Not from my standpoint.

FRED BAKER: Then I guess we're done and we'll talk with you guys next month.  
Thanks a bunch.

MARIO ALEMAN: Thanks everyone, this call has been adjourned.

**[END OF TRANSCRIPTION]**