
FRED BAKER: [inaudible] proposed and Geoff agreed to last week was that all talk about his lab and experiments that he's doing, and Geoff talked about his lab or his experiments and things he's doing and we kind of generally discussed questions that people would like to explore or would like to be explored.

Is there anything that anyone wants to add to that agenda?

MOHIT BATRA: Hi, Fred. This is Mohit from New Delhi.

FRED BAKER: Yeah?

MOHIT BATRA: Okay. I [inaudible] some of my suggestions some time back. I guess everyone had a chance to look at them. So, if anyone has any comments on any of my suggestions, we can discuss those.

FRED BAKER: That sounds good.

MOHIT BATRA: Okay, thanks.

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.

FRED BAKER: Paul, I'll tell you what. Let me turn the meeting over to you. You can talk about your [inaudible].

PAUL HOFFMAN: Okay. That sounds fine. Some of the things that I will talk about actually would hit a few of Mohit's questions, although more of those would be in the second part, if Geoff shows up or I can cover them there as well.

At this point, given the question ... And again, we're totally open to folks in the work party or just in the caucus, since the work party is a little bit of a loose thing here, suggesting new things, ways to interpret or extensions to what RSSAC asked us to do.

But it seems that the tasks for us, the research tasks, are split into two broad categories. Tests for what do the current resolvers do – I'm sorry, the current resolver software do – in particular environments, with the main one being how do they prime, how do they get attracted to particular authoritative servers, especially root servers. So, that's one section which is, again, looking at software implementations and that would be done in the lab.

Then, the second section is what do we see in the wild and what can we do to ... How can we do better testing to make some of the things we see in the wild more testable and things like that? That would be many of the things that Geoff Huston has already been doing with his APNIC work. But we'll talk about those later.

So, let me focus on a test bed. Let me just see. I might be able to share the document. Lord knows if this is going to work. Give me a moment

here. This might actually work. Resolver test bed. Let's see if PNG's work. Yes. Okay.

This is the design that I'm working with currently. The host box is a control host, which would be any old laptop or desktop computer that can run a VM system. I'm currently using virtual box, even though it's not quite as good as VM Ware, simply because since we want these tests to be reproducible by anyone, we don't want to force people to buy software just to do it and virtual box seems, so far, to be working well enough.

So, on the control host, which just happens to be my Mac laptop but it would work on many box that can run Virtual Box, I have set up three virtual machines or the plan is – and again, all of this is with input from the group. That would be great. Three different virtual machines. On the left, would be authoritative name servers which we would use essentially for doing the testing. That is that when the resolver that is being tested sends queries, they're going to hit something in our test bed, so that we can very carefully control the timing and things like that. So, it'll probably just run BIND as an authoritative. We might even be able to test it if we think it's important with some other authoritative servers.

On the right-hand side would be the resolvers that we are testing and that currently is planned to be any resolver that we can run in a virtual machine. So, that's all of the popular resolver software. But, if somebody, for example, has something where they want to test something that cannot be run in a VM, such as a resolver that's in a piece of hardware or something like that, as long as it can connect to

the test bed with these two interfaces, it should work okay. We probably won't get to that testing until later.

Then, between the servers and the resolvers is a gateway and the gateway is going to be where we are both going to be doing all the packet captures, so that we can see the ...

But, also the gateway will be doing things like slowing down queries – I'm sorry, slowing down responses from someone's servers, so that we can test attraction to the servers. What this means is we're going to have to have our own root, but that's fairly easy in this controlled environment. All we need to be able to do is to push in a different root [ends] file and then we also would need to push in the new root zone with the test bed in it would also need to be signed, which DNSSEC signed, so we would also have to be able to push a trust anchor.

All of that is fairly easy and we have experience with this in my group at ICANN because this is pretty much how we set up the previous tests we were doing for the KSK rollover, so I sort of stole a bunch of the [inaudible].

So, that's basically what I'm thinking now. I'd be interested if people had questions or thoughts on this.

FRED BAKER: Well, I have a vision of—

WES HARDAKER: I think it looks good overall, Paul. Go ahead, Fred, sorry.

FRED BAKER: Oh, I was just imagining this being documented as how to build with] an alternate root.

PAUL HOFFMAN: Yes, although the alternate root part would be pretty trivial for anybody. It's just signed with something different and push out things that are different. The nice thing is if someone followed the recipe for this exactly, it would not help them at all because the [serve net] network and the resonate network are both on private address spaces, so if they actually used exactly the IP addresses we used, no one would be able to see it, which is exactly what we want here because we want the DNS packet traces that are being kept on gateway VM to be absolutely as clean as possible with no outside interference, which is why, for example, the servers VM does not have a direct connection to the Internet itself. It has to go through the gateway VM so that we will be even able to see, for example, when we start up the BIND authoritative servers with this alternate route, what kind of queries do they send out to the Internet. We will know exactly because we'll be capturing those as well. Wes, you had a question?

WES HARDAKER: : Yeah. Actually, there's already stated documentation for how to [stand up] a root server. It's called Yeti. And they've done it. Anybody that wants to do an alternate root can already do that today. It's not terribly complex. Overall, Paul, I think this looks like a good architecture. I'm looking forward to it being up and running.

There is some technology for doing delays and stuff like that, that already exists, dating back to [Emulive] and then followed on into [TOR] and stuff like that. I assume you don't need to write this yourself, right?

PAUL HOFFMAN:

Right. I was going to be using IP tables and TC since we're using IP tables on the gateway VM anyways and it has TC on it. I've been told that's actually quite easy to do, that in IP tables, you simply set up a bunch of labels, one per address. So, the servers VM will run all 26 addresses, 13 IPv4 and 13 IPv6. So, we can even do tests, for example, that say how does the priming and the repriming work if v6 is significantly faster than v4, for example, for all of the servers and such like that.

So, yes. I will probably get to that sometime early ... So, I can easily do what we've got here, have this publishable certainly by around beginning of January. I hope to also at that point have some of the delay stuff built in. And if so ... If I'm having problems with that, I would come to the list, but I'm hoping TC will work. And Wes, can you send me pointers to the other ones, just in case those are easier to use?

WES HARDAKER:

Yes, I can. And TC should certainly work as well. In fact, that might be what they're using internally.

PAUL HOFFMAN:

With all of these, it gets a little bit [fiddly] with setting them up.

WES HARDAKER: Yeah. [Traffic shaping] in Linux requires some reading, just to properly understand it. one clarification question. So, the servers VM, that's a single VM running 13 on one box or is that 13 different VMs? No, which?

PAUL HOFFMAN: No, that's a single VM running 13 on one box.

WES HARDAKER: On different ports?

PAUL HOFFMAN: No, single port. Or actually 26 addresses.

WES HARDAKER: Okay, so 26 addresses and it's responding to whatever address it receives on.

PAUL HOFFMAN: Correct. I'm sorry. So, I guess I haven't actually worked this one out yet, but I assume I could probably run 26 instances of BIND, each one listening to one address. I also think I could run it where one is listening to all 26 addresses, since the only thing we're really doing that would change between addresses is delays and such and that would be being done on the gateway. Do you have opinions on that?

WES HARDAKER: No, I'm just trying to think through it. Is there any ... Are we losing anything by using a single VM as opposed to 13 VMs?

PAUL HOFFMAN: Right. That's actually—

WES HARDAKER: And I can't think of anything, so ...

UNIDENTIFIED MALE: But, if we are, it certainly wouldn't be that hard to do this. I just sort of suspect that when 13 identical-looking interfaces on Linux will cause mysterious problems.

WES HARDAKER: [inaudible] with aliasing [inaudible].

UNIDENTIFIED MALE: I know you can. I'm just ...

WES HARDAKER: What you're doing should work fine. I can't think of any downsides, but my academic mind is like, "What's wrong with it?"

UNIDENTIFIED MALE: No, no, that's great, because again, since this is all running on one control host, we could certainly extend it in any ways that we need to. But my thought was a single [inaudible] with either 26 instances of BIND or one instance of BIND looking at 26 addresses, as long as all the fussing is happening in the gateway.

And again, [serve net and resnet] in this diagram are on private address spaces and they're on separate private address spaces, just so the servers can't talk directly to the resolvers for some weird reason. Oh, and then, just to be clear, the gateway VM is running [anat].

FRED BAKER: So, you're using the IPv6 address – are you using a [ulader] or something?

PAUL HOFFMAN: Haven't gotten that far yet. My IPv6 experience pretty much is I have Groundhog Day. I need to relearn it every time I have to do any routing with it. So, I might come and ask you, Fred.

FRED BAKER: Okay.

PAUL HOFFMAN: So, we're using a bunch of design that our group already has. I don't think it will be any problem for me to have this ready in January where not only that it works, but it's also the instructions will be good enough

for y'all to be able to try it out yourself, as long as you have virtual thoughts. And pretty much it should be all you need is virtual box and a recent ISO for [inaudible]. I believe that will be enough to at least get us started.

Any other questions on this? I know that you're just looking at diagram. You will want to see tests at some point. I'll try to push some out.

Okay, has Geoff come on yet or not?

FRED BAKER: I haven't seen him.

PAUL HOFFMAN: Okay. Geoff, are you on under somebody else's name? Okay. Well, let me do a general talk on second set that – actually, Joe, can you [inaudible]. Great, thanks.

Actually, [inaudible], can you speak for Geoff since you are part of that team? Which would mean you have to unmute, that I just asked you to mute.

UNIDENTIFIED MALE: Hi. Can you hear me now? Well, what were you asking me to talk about?

PAUL HOFFMAN: Basically, his design for how APNIC [inaudible] does [inaudible] of various [inaudible] resolver, particularly on IPv6 capability, [inaudible].

[JOAO DAMAS]: Okay. I can speak a bit to that since I actually run the infrastructure. So, everyone knows we use experiments that run [inaudible] script code in [quiet] machines throughout the world and the way we inject this javascript into the [inaudible] machines is in [inaudible]. In this case, the double-click Google ad platform, right? So, of course, what we can do in javascript inside the browser is limit it, which is a good thing. But one of the things you can do is fetch a URL, so using the fact that you can send the browser to a URL of your choice, what we have is we designed names that [inaudible] using uniquely [inaudible] set of five, six, seven, whether or not we [inaudible] minimum time.

Each of these names, when the client has to resolve it [inaudible] others to make a connection has to go through a certain number of DNS [inaudible]. So, for instance, we have experience where the name is only possible to resolve to an [inaudible] and you can only [inaudible] using before. We have a separate experiment that runs in the same batch, where the same is true but only using the six, so you can only fetch [inaudible] response. Eventually, the [inaudible] which is the [inaudible] for that name.

We have experiments where the response [inaudible], so we then observe who asks for [inaudible] for both, in what [inaudible], with what [inaudible] with respect to each other. Then, we see on the website connection what was the decision of which one they would pick to

finalize the connection or whether they [inaudible]. Sometimes, [inaudible] over to one person and then the second one.

That's how we do the measurements [inaudible] together with preferences and delays of who is faster, which [inaudible] preferred by the client.

So, the interesting part of course in this is that there is a DNS component, which you usually are exercising the resolver that the user uses, not the [inaudible] on the web browser or the app, the mobile app, that is displaying the ad. But actually the resolver that people are using. Some people are using the ISP resolver and some people are using [inaudible] Internet, of course. So, we get the also good mapping based on the IP address of the client which we see both when the client comes to us asking for the [inaudible] of the experiment and then afterwards when they do each of the web [connections]. So, we can place them in an [inaudible] system as well as [inaudible] in an economy.

Then, we see which resolvers that client was using, basically because we can associate through the uniqueness in the name, we can associate a client with a set of resolvers that asks that question on behalf of the client.

Interestingly enough, of course, you also have to take into account that there are people that ask for the questions that the client was asking in the beginning but are not the ISP resolvers. So, people who, for whatever reason – some reasons are good, some are debatable – [inaudible] going through the wire which on which they [inaudible].

So, we see this thing of [inaudible]. So, sometimes we have [inaudible] to address the specific aspects of the resolver behavior. One of them that is long-running is whether they have validating or non-validating resolvers and we do that, for instance, by providing two experiments, one where we provide [inaudible] with a valid DNS signature and a [inaudible] experiment where we provide – we break the signature, so it's the same answer, but the signature that's attached is broken. So, if you are actually validating, you shouldn't follow the second [inaudible] but you should follow the first. If you are, for instance, following both, it probably means [inaudible] information, you are not actually doing anything with it.

On the website, the amount of things you can do is usually more limited and we lose that as a final closure to the experiment of how many people actually manage to go through the whole process of actually achieving the correction.

These last [inaudible] have a little bit more knowledge to [inaudible] than DNS because the DNS basically is a synchronous system. Once the client asks the resolver to do something for it, then even if the client does a [inaudible] the connection is terminated, [inaudible] DNS part, the resolver will still carry on doing whatever it was asked to do, even if [inaudible] back to the initial plan, the plan is no longer there.

The website, on the other hand, requires that the client stays on that activity for the duration. Most of the time, we do. It's [inaudible] most of the time. But there are places, there are combinations, there are time outs, there are [inaudible] that lead some experiments and some ISPs

and some clients to take a [inaudible] time, so you have to take account those because the [inaudible] might not be asking a single ...

But the [inaudible] is very good [quantity]. The good [inaudible] this way of [inaudible] client is that you reach millions of them and distribute it everywhere. So, that's a quick overview.

PAUL HOFFMAN:

Great, thanks. So, [inaudible], do you have a ... Since there's only a few of us on this call, do you have a write-up, basically, of what you just told us somewhere, where we can then share it with the work party?

[JOAO DAMAS]:

The answer is yes. There was a presentation and some [inaudible] a couple of years, three maybe ago. I don't recall the URL right now of the type of [inaudible] but I can make a note and send it to the caucus list maybe.

PAUL HOFFMAN:

Sure. As long as it's [inaudible] related to this [inaudible] specifically would be one of the things that we're trying to do is have the measurements that we take be reproducible by other researchers. I remember the presentation you're talking about. It was a little bit light on how you do the [inaudible], which I guess means how you, [Joao], does the tooling.

Is there anything or are there any plans to expose that a bit better, so that other researchers, even if they don't have access to what the

Google Ads could ... If they can get, inject all the script into hundreds of thousands or millions of users in other ways, run the code or run their code and then you can get comparative results. Is that possible?

UNIDENTIFIED MALE:

I would think so. I mean, the [injection] platform is whatever. You use your choice. If you are in China, you might use [inaudible] or whatever, Google, Bing, pick your choice. They all work the same, basically. The java script is just [inaudible] what we have is basically a system is where the first thing the java script does is connect to one of our web servers and fetch a list of URLs, which is the then current list of experiments that we want, and by list of experiments, I mean the names. Then, the trick is always on the DNS software that we use, [inaudible] service. That was code that was originally written by [Ray Vallis] [inaudible] and it's on GitHub.

Of course, for each experiment, we have to do different things, so we change the code base for that name server fairly frequently, so it's not a code base that will be worth anything to anyone else because it's specific to the experiment. But, the base idea of [inaudible] DNS, [inaudible] DNS – basically it's the combination of [inaudible] DNS library, and basically what it provides is it's a platform that is easy, relatively easy at least, to modify, to run those experiments. So, it's a DNS server that can be made to lie in whichever way you want for your purposes and it is easy to [modify] under that. That's already on GitHub. If you look for [Ray Vallis's] GitHub [inaudible], you'll find it there.

After that, it's just processing the [locks] and there is no mystery there, really. It's just [inaudible]. Geoff likes [inaudible]. We just [inaudible]. It's not a big deal. The secret sauce is on GitHub. It's a [inaudible] of that code I mentioned.

PAUL HOFFMAN:

Great. So, specifics for your test would not be of interest to anyone. I disagree. I think that it would actually be helpful for somebody – and I don't think this will be me, but it might be – who wants to, again, reproduce your results, essentially with a different customer data set, to actually be able to start from your setup, obviously on their own servers, and then be able to extend from there.

For those of you who haven't looked at – I always call it EVL DNS. I don't [mind] the extra syllable to make it not seem so evil. For those of you who haven't looked at EVL DNS, it basically lets you write C code to do the specific changes that you want.

So, us seeing some of those, especially for the tests that the work party are most interested in, which I think at this point, the top ones would be is the resolver IPv6 capable and is it doing DNSSEC validation would be useful to us.

Again, I have not seen [you folks] publish them and I don't know whether you are allowed to with APNIC or whatever, but it certainly would be helpful to us to see at least a detailed description of the design of your test, if not the actual code.

UNIDENTIFIED MALE:

Well, certainly you can write. I don't think [Jack] would have any problems. It sounds like the last thing [inaudible] a little more detail. If you look at the EVL DNS – I'm sorry, I'll keep calling it [evil] DNS because that's what [Ray] wants it to be called, so I will honor the original offer. If you look at the code, you will see that the base code, basically you give it a set of zone files and something like a wild card answering machine. So, it will answer [inaudible] on that zone that you configure it to do based on the content of that zone file. If you put the [inaudible], they provide the [inaudible]. If you put the [inaudible], they provide the [inaudible]. If you put both, it provides both and so on.

The breaking of the DNS signatures is trivial. When we tried to use DNSSEC properly, it used DNS to sign the records and these [inaudible] in the code that's on GitHub.

The other thing that is done by [inaudible] on the real signature. So, basically, you flip a few [inaudible] and it breaks the signature as it should and that's it. So yeah, we can write more details, but the code is actually there for those basic experiment.

It's only one [inaudible]. You said C. It's actually C++.

PAUL HOFFMAN:

Thank you. Right. [inaudible] as carefully as I should. Good. If we can get back from you folks also a more detailed overview of how you want to take care of tasks and such, I think that would be helpful to those who want to reproduce, or folks who want – even if they don't want to reproduce your work, if they say, "Gosh, I've seen this result from you, could you also add in this little thing?" since like you say, you were

running these pretty constantly and you can modify which tests are getting done [inaudible]. That might be useful for the work party as well because we may, if we're not running – if we're not reproducing what you're doing, we can at least narrow down some of the things that might be of interest to the folks at RSSAC.

UNIDENTIFIED MALE: And we would be happy to take suggestions of new aspects of experiments of running them and interact that way. I don't think Geoff will have a problem with that. I'll certainly have to check with him.

PAUL HOFFMAN: Okay. That sounds good.

WES HARDAKER: So, [inaudible] have any of that.

PAUL HOFFMAN: Do other folks have questions for Joao? Okay, hearing none, I think, Fred, that's the two parts of our experiment of how we will set up test beds and look at experiments and such. Over to you.

FRED BAKER: Okay. So now, the other thing that you mentioned in your proposed agenda was does anybody have any other questions? Mohit, you sent a note and I just went looking for it, and for whatever reason, I don't see it. Do you have any questions remaining that we haven't touched on?

MOHIT BATRA: Not really right now, but I'll go through the minutes of this meeting again and if I have some more questions, I'll write down to the mailing list. But, in the meantime, the comments which I made, if somebody has any comments on those probably, you can discuss.

PAUL HOFFMAN: So, Mohit, I can respond to some of those now since we've covered both sides of the intended testing. And Fred, if you wanted to look, Mohit's message was sent to just the work party and it was on November 12th. The first question was that we should utilize a combination of commonly used DNS server software, recursive and authoritative. That's certainly what we will do in the test bed, in my test bed, as you've seen.

I wasn't really thinking of the authoritative software side as much, although that server VM could indeed have ... We can run experiments where that we can see, for example, what if all the root servers were running NSD instead of BIND and such like that.

But, for the resolvers, where we want to test as many resolver software as we can, that resolvers VM, it is my intention to test as well as I've done some of the other stuff, which will be literally dozens of versions of BIND, Unbound, [knot resolver], Power DNS, and such like that.

Does that sort of answer your concern?

MOHIT BATRA:

Yes, certainly. As I also wrote, Google public DNS is one of the most widely used resolvers. And I'm not sure whether we can assimilate the exact Linux configuration for that, but I mean, how do we actually assimilate that Google public DNS software? Because it's most widely used.

PAUL HOFFMAN:

So yes, I've talked to folks at Google over time and I'm sure the answer is no, that we cannot, because their setup is proprietary and it is constantly under change, and I've been told that, for example, even if we had send out queries from 100 places around the world at a given time that are all going through 8888 to authoritative servers that we control, so we can see how it happens, we will actually see different patterns. That is, it is not at any given moment guaranteed to be a single universal system. But we can certainly do that. I mean, I understand your concern that since this is by far the most popular resolver, shouldn't we know more about it? We could set up a set of queries that all go to authoritative name servers that we run, very similar to what Joao was just describing for the way they do things, and get some results from that.

We might want to do that for the top five open resolvers, so that we're not just looking at Google, but to look at quad nine, quad one, what used to be VeriSign, things like that. I think that would be reasonable, and especially if you were interested in organizing some of that, that would be great.

FRED BAKER:

You had your hand raised. Do you want to get in here?

UNIDENTIFIED MALE:

I did have a question to Paul and actually Wes, but I just saw Wes was dropping off. It's just a question we can answer later. I remember Wes was, during the IETF meeting, discussing that he also wants to have a more analytical model – maybe not the correct wording, but kind of describing how resolvers behave. So, of course, the open source developers can help with, and maybe in this model, if he can get in contact with Google [inaudible] [inaudible] already answered. They don't know either how it works, but at least get some generic behavior how their pool of resolvers behave, especially towards the root servers. But, actually, it's more a question to Wes if it's relevant or important for him.

UNIDENTIFIED MALE:

Well, I think it's actually important. It would be a useful thing, assuming we have time, because remember, we do have a one-year deadline on this work party. If we can fit it into the one-year deadline, I think that what he brought up, it would be useful in general, which is if we can ... If some of us can specify what we expect or what we say is expected resolver behavior, we should be able to test for that, both in the test bed and out in the public.

And for those of us who have seen Geoff Huston's presentations on some of the results he gets when he runs it, and of course everything is all over the map, but [inaudible] resolvers. I think that would be useful, because again, we, at least for any of the open resolvers, we can set up

tests for them to see are they doing what we expect, regardless of what we think their design is?

And I didn't want to indicate that Google refuses to say. I think they're hesitant to say, because then people will be attached to what they say and they change their stuff over time. Just as [inaudible] pieces of – many of the resolver vendors do.

UNIDENTIFIED MALE: I understand their problem, but [inaudible]. Yeah.

PAUL HOFFMAN: Yeah. But I think that would be a reasonable result that RSSAC would like is if we came up with a list of expected resolver behavior and then we could test it against what we have in our various test beds and see what we get. And that would sort of go along I think with what Mohit was asking about the Google public DNS is what can we learn about it?

MOHIT BATRA: The [inaudible] second and third point, it might be beneficial if we create a repository for study material, probably some research papers which are best mentioned in an earlier meeting.

PAUL HOFFMAN: Do you have a list that ... Like if I came to you right now and asked you for a list, do you have even the start of such a list of those types of papers?

MOHIT BATRA: Papers I don't have, but certainly I can try to find a doc.

PAUL HOFFMAN: Okay, because we can easily set up a Google doc page that is shared by everybody that people can start working on. That can be sort of a parallel project. I think that would be a useful outcome for everybody.

MOHIT BATRA: Okay. I'll try to collect this information.

PAUL HOFFMAN: Great. Thank you. Fred, I think that covers all of Mohit's issues from the message. Of course, more might come in later as we discuss them on the list, but I think that covers it all.

FRED BAKER: Okay. Thanks much. And Mohit, I have looked around for the e-mail that you send on November 12th and I don't have it. I don't know why. Could you re-forward it to me, just so that I have it?

MOHIT BATRA: Yes, I can do it right now.

FRED BAKER: Okay, thank you. So, we've got another 13 minutes to the top of the hour. Does anybody else have any questions you want to raise or other topics at this point?

Hearing none, I'm going to give you the rest of the hour back. Well, there is one question that we should talk about. When do you want to talk again? Do you want to have another call in January? Is that too soon?

For people in the US, most of us are going to be kind of out of commission for the next month.

PAUL HOFFMAN: By the middle of the month, I will certainly – I can commit to having a good shot at the test bed that I showed the diagram from today. I mean, I would certainly send it out before such a call.

I would say if we schedule for, say, five-ish weeks from now, we would at least be able to talk about that and maybe Joao will have some of the stuff from Geoff's side as well by then.

FRED BAKER: So, five weeks from now would be about the middle of January.

PAUL HOFFMAN: Yes. So, it's [inaudible] today is the Thursday, if we did a Thursday again, the [inaudible].

FRED BAKER: Okay. Well, what I'll do, then ... I'll do it or maybe Mario will do it. Boy, there is a weird echo on this. We'll send out a Doodle poll and just make sure people are available. Basically, targeting that week. Oh, the echo went away. Did everybody go with it?

PAUL HOFFMAN: No, Joao muted.

FRED BAKER: Okay. We'll send out a Doodle poll and make sure that everybody can agree on a time. Thanks a lot. We'll go from there. But let's make this meeting adjourned.

PAUL HOFFMAN: Okay. Thanks, everybody.

UNIDENTIFIED MALE: Thank you, everyone. This meeting has been adjourned.

PAUL HOFFMAN: Thank you.

UNIDENTIFIED MALE: Thanks, everyone.

[END OF TRANSCRIPTION]