MARRAKECH – RSSAC Work Session: Modern Resolver Behavior
Wednesday, June 26, 2019 – 13:30 to 15:00 WET
ICANN65 | Marrakech, Morocco

FRED BAKER: Okay. So, we are now in the RSSAC Caucus Work Party meeting related to understanding DNS resolvers. Paul is actually the man of the hour with respect to it and he's online, so we're mostly going to hear from him, I believe. We have the agenda ahead of us or on the screen ahead of us and what, Paul, you start out by looking at the statement of work and then the work that you've done. Let me turn the agenda control over to you, Paul.

PAUL HOFFMAN: Sure, that sounds good. Although, don't we need to do a roll call first?

FRED BAKER: Oh yeah, okay. Who in the world is here anyway? So let's start from Wes.

WES HARDAKER: Wes Hardaker, USC-ISI.

RYAN STEPHENSON: Ryan Stephenson, DISA.

PAUL MUCHENE: Paul Muchene, RSSAC caucus.

SHINTA SATO: Shinta Sato, JPRS RSSAC caucus

LARS-JOHAN LIMAN: Lars-Johan Liman, Netnod.

HIRO HOTTA: Hiro Hotta, WIDE and JPRS.

MICHAEL CASADEVALL: Michael Casadevall, ICANN Fellow. I did a lot of research on recursive resolvers.

OZAN SAHIN: Ozan Sahin, ICANN support staff.

ANDREW MCCONACHIE: Andrew McConachie, ICANN support staff.

FRED BAKER: Fred Baker, RSSAC co-Chair.

BRAD VERD: Brad Verd, RSSAC co-Chair.

HOLLY RAICHE: Holly Raiche, ALAC.

JEFF OSBORN: Jeff Osborn, ISC.

RUSS MUNDY: Russ Mundy, RSSAC caucus and SSAC liaison to RSSAC.

DUANE WESSELS: Duane Wessels, from Verisign as the root zone maintainer and liaison to RSSAC.

DANIEL MIGAULT: Daniel Migault, Internet Architecture Board, liaison to RSSAC.

KARL REUSS: Karl Reuss, University of Maryland, root server operator.

KENNETH RENARD: Ken Renard, ARL, RSSAC.

UNIDENTIFIED MALE: [Inaudible], RSSAC caucus member.

JOY CHAN: Joy Chan, TWNIC.

UNIDENTIFIED MALE: [Inaudible].

RICK WILHELM: Rick Wilhelm, Verisign.

FRED BAKER: Okay. So, Paul, now I'm giving you agenda control.

PAUL HOFFMAN: Okay. I thought that Ozan would be reading the folks who are remote participants because there's a bunch of us.

OZAN SAHIN: Yes. Thanks for reminding, Paul. In the Zoom room we have Champika, also Holly Raiche who is also in the physical room. We have Kazunori Fujiwara, Keith Bluestein, and we have Hiro Hotta who is already in the physical room. We have Rick Wilhelm from Verisign and we have [inaudible] and Ryan Stephenson.

PAUL HOFFMAN: Okay, great. As Fred had said, the agenda for today's work party is a little bit different than normal. In the last few work party meetings, we've mostly been talking about the code base for DNS resolvers and the test bed that we're using to look at that particularly for things like prime inquiries and such. And we've had a little bit of discussion on the mailing list about how the actual resolvers on the Internet work which was the second part of what RSSAC asked this work party to deal with.

So, today I wanted to focus pretty much strictly on that so that we can figure out if there is still interest in moving this forward, which would take a bit of commitment from work party members. And if not, we can sort of wrap it up for now, but I thought we should take another strong hit at that. I apologize for the lateness of getting this message out to the caucus and the work party, and I know that you folks were all at lunch. I'm sitting here watching the sunrise.

The statement of work had four parts. Part one and part four are on the screen there, which is basically to say to analyze network traffic for behavior, to understand how resolvers interact with authoritative servers in general and RSS specifically, in terms of preferred root server selection. And then analyze them to understand just how they interact with the RSS in general. So,

we're not going to try to hit number one at all in this work party. I believe simply because trying to figure out about how resolvers do preferred root server selection by looking at traffic is extremely difficult. We get all sorts of conflicting data when we look at DILO data and such like that, and I just don't think that it is terribly useful. However, as the research APNIC has done, there are actually interesting things we can do to see how DNS resolver system interact with the root server system in general.

For example, APNIC has done a number of studies, some of which are ongoing which come out with results such as what percentage of the resolvers out there actually can do IPv6 successfully? What percentage of resolvers are doing DNSSEC validation? And of course, there's always problems with any of these measurements but some of those are in fact of interest to the root server operators, especially some of the data that we're seeing from APNIC right now which could be of great value, indicates that the number of users behind where the resolvers are doing DNSSEC validation maybe going down, not up. That may be a temporary thing, it may be a measurement error, but things like that could be of great value.

So if you can scroll down a little bit, Ozan? In the message I sent – and again, I apologize for having sent the message out just barely more than an hour ago – this is a summary of how APNIC – the resolver tests work. Geoff Huston from APNIC has actually

sent more detailed information on this to the work party. But in general, the way that they do this is that they have some JavaScript code which causes DNS lookups for unique names. They have a mechanism to get that JavaScript in front of a large number of users, and when I say "get in front of," that means "And execute it."

Then in order to do the analysis part, they have custom authoritative servers, they get queries and give customized answers and maybe do other things. Now, as a note, RSSAC asked the work party to be looking at resolver behavior. Some of the things APNIC currently measures goes beyond that and so for example, in their IPv6 testing not only do they come out with numbers that say approximately what percent of resolvers can handle IPv6, but they also look at whether the users behind those resolvers can handle IPv6. That's probably out of scope for what RSSAC wants us to do, but just to be clear, using the setup like APNIC has, you can go a bit further and find out more about the users.

In our earlier brief discussion in the work party, we realized creating JavaScript that does this is fairly straightforward and also setting up the authoritative servers takes a bit of configuration, but this is not rocket science. Where things would be interesting if the work party wanted to pursue this is actually

how do we get the JavaScript in front of a large number of users so that we can get the kind of results that we want?

The way APNIC does this, as many people know, in fact, people often just refer to this as the Google Ad network, is that they have an agreement with Google, and Google gives them free or virtually free advertising, but they don't have a whole lot of control over the advertising. Now, Google has done a fairly good job for them over time of putting the ads that have the JavaScript in them in front of a bunch of different users, and so that APNIC can run multiple long-term test and such like that. But it's not an agreement where Google will say to them, "Here's exactly the kind of users who will see this ad, therefore your JavaScript.

Now if you scroll down more to that following list. Yeah, keep going. Yeah, okay. So, we can stop there. So, in the work party there was a desire to if we are going to do this work to complement what APNIC is doing, and one way to complement that would be to use some sort of different delivery mechanism. That is that going and doing what APNIC is doing with the same delivery mechanism – go into to Google as a separate group is probably not that interesting just because – in fact, if we can look at all the parts of what APNIC is doing, it's much easier for them to do that for us. But there's been many questions about how APNIC's delivery mechanism is really representative. For

example, Google ads are not shown universally throughout the Internet environment and, in fact, APNIC has discovered that different kinds of ads show up in front of other people.

So, looking at other delivery mechanisms, people have mentioned, "Well, there's other ad networks." We could also put the JavaScript just on webpages. Now, putting it just on a webpage means we have to induce a fair number of users to go to the webpage, but there are some web properties that are pretty much international in scope. And I'm not only speaking of social media, because actually social media has many of the same limitations that the Google ad network has, that it doesn't get everywhere. Some countries block social media. Some countries block social media at different times of the year, things like that. But there are businesses that have users throughout the world and I think of some of the physical delivery systems such as UPS, DHL, and such like that, where if one of them allowed us to put the JavaScript on their site for a while and they spread that across all of their international systems, that that could be considered international. Large companies like Salesforce also have a very large worldwide footprint. So, that would be a potential additional delivery mechanism.

The third one I list here are games, things that are not websites but have an international scope. And as someone pointed out to me, there are actually more users on a day-to-day basis of Candy

Crush than there are in fact users of the top Alexa websites starting at number 10 or so. So, possibly getting together with a game company that will allow us to do this. Because the games of course also have a lot of code in them that could easily be used to generate JavaScript.

So, each of these delivery mechanisms has positives and negatives and I have not exhaustively list them here, although I think that that would be most valuable to discuss in the work party. Different advertising network will hit different people. Some advertising networks are really based around text-based searching, which would hit a different population than those who are looking at videos. And even though YouTube is the most popular video source, it's by far not the only way to be getting videos. For web properties, if we went to business web properties, we would then not be hitting youth as much. If we go with games, we're probably going to hit youth and not as many of us old farts. Somewhere maybe a combination of them.

If you can scroll to the last bit of the message, Ozan, and then we can start this sort of freeform discussion. Some of the questions that come up and that Geoff has or APNIC has had to grapple with in delivery of the JavaScript is, should users be prevented from contributing to a measurement more than once? That is, if I click on two different pages and get the same Google ad, should I be participating twice? Or should I be prevented from

participating even if I've moved networks? I have my phone on here at the house, which means I have one ISP and therefore one resolver, I then take it out and since it's on my local Wi-Fi, I then go into my car and drive away, and therefore I'm on my mobile provider's resolver. Then I take the phone and I go to a place where I go off, and therefore I'm on yet another Wi-Fi network. So, should we be tracking the user multiple times or not? That would have a fairly large effect on the measurements that we make.

There's also the question of, do we want to try to drill down and find more about the user as a way of figuring out how users use resolvers? That's a very open question because the more information that you try to get about a user, the closer you get to personally identifiable information. And even if you do a fair amount of scrubbing, some of these can be fairly invasive. As APNIC has said before, they are quite restricted in their current measurements by Google about the kind of level of information they can get from a user in this free network environment. And quite frankly, even if you were paying for your ads, mostly ad networks do prevent advertisers from scrubbing too much information about individual users.

Then this sort of comes down to once we have some values and we know a little bit, do we care about what kind of users these are – home users versus office users, that is, which can often be

considered small landline? Landlines with a small number of individual devices versus office. Or even when they are doing it. A daytime user will have a different profile than a nighttime user, assuming that most people are "working" during the day or in a workplace, and at night they are at "home."

In summary, some of these will depend on what RSSAC wants us to measure. As I said earlier and as Geoff has emphasized, that the APNIC information, that they are publishing goes definitely towards measuring users and especially number of users behind certain resolvers, so that they can talk about the most popular resolvers, the resolvers with the most users and such. That may or may not be of interest to RSSAC, so sort of the two ends of the spectrum is that you want to measure as many users as you can filtered through their resolvers possibly at a certain time of day or not, or simply that RSSAC only cares about how do the resolvers themselves interact with the root server system.

And so, during a test run any individual who's behind any resolver, that's good enough. We don't need to get a bunch. We don't need to measure them. We're really trying to get in the scenario just at least one hit per active resolver after.

With that, I would like to open it up to anybody in the room and especially to the RSSAC folks who created this statement of work, to sort of give us a little bit of direction and then see if

there's interest in pursuing this. And since I'm not in the room, Fred, I'll ask you to take over and point at folks in the room and to be looking over at Ozan in case there are raised hands from the remote participants.

FRED BAKER: Sure. I'm willing to do that. Okay, so does anybody have questions or comments they'd like to make at this point?

RUSS MUNDY: Yeah. This is Russ. What I thought – and I have not been following the work party really closely – but I thought the focus of the effort was to get characterization, if possible, of the different types of resolvers out there. I guess if I'm understanding what you were just describing, Paul, that this is still the basic core objective of the work party but you're looking for more guidance as to what kind of characteristics are desired?

PAUL HOFFMAN: Yes. So, when you say type of resolvers, I'm not sure if you mean brand and version of resolvers or their connection characteristics or their configuration characteristics. For example, APNIC's – the two measurements that I report – and they have many of them – are: is this resolver connected over IPv6? And I believe the APNIC has actually even been studying,

and how good is that connection? A configuration parameter would be, is this resolver validating? That is DNSSEC validating. If those are the kinds of things that RSSAC wants from us, we can do those kinds of measurements. Maybe they want other things from us. And one of those other things could be, how popular is each resolver?

MICHAEL CASADEVALL: Michael Casadevall, ICANN Fellow. I hope I'm not [aligned] as an observer here but I've been actually building a rather extensive toolkit for measuring the behavior of recursive resolvers and understanding the behavior such if they're properly validating DNSSEC and such. I've done some work also in seeing how this work can be integrated into tools such OONI Probe and as I am also [inaudible] developer, it may be possible to convince those communities to shift DNS test tools from the default distribution as both those distributions collect telemetry information and maybe willing to assist with this. Given the right context and information, I can make some inquiries if this working group is interested or discuss further on the work and research I've been doing along these lines.

FRED BAKER: I think you just said that among other things, you could increase the probability of a DNS record being validated, is that correct?

MICHAEL CASADEVALL: My project is called DNS capture. It was basically designed to catch DNS resolvers lying in the act, partially motivated by the fact that my home ISP does this quite a lot. For example, it says everything is DNSSEC validated and captures any NXDOMAIN records and sends them to the wonderful advertising sites. This was a larger built out of a project to understand the behavior of recursive resolvers. I presented the initial design at Internet Freedom Festival and applied for funding from the Open Technology Fund. I have a working prototype of this code in place and could modify it fairly easy to include RSSAC requirements and then work on helping it get deployed through various organizations that do this sort of data collection such as OONI as well as trying to integrate it into various operating systems as an open measurement which would get a pretty good cross-section of users.

FRED BAKER: Okay. That actually does sound interesting. I'm going to suggest though that we probably need to have a more in-depth discussion about what you're doing and that kind of thing. So, I'd be happy to talk with you later. I'd encourage you to talk on the – you're a member of the caucus, right? No, you're not? Okay.

MICHAEL CASADEVALL: I saw the description, the work, and saw the two lined up, so I sat at the table and hoped to – I wasn't being out of line.

FRED BAKER: So, as a bureaucratic hurdle, if you want to do something with DNS, we'd like you to join the caucus.

PAUL HOFFMAN: Fred, could you get closer to the mic, please? You're barely audible.

FRED BAKER: You mean I have to actually use the mic in this? As a bureaucratic hurdle, working in this context, we'd like you to join the caucus because that's the set of people that help us to sort things like this out. And if you go look on the Web, you'll find a form you can send in and we'll take care of that. But let me encourage you to talk with Paul, the guy on the other end of this connection about this stuff that he is doing and how yours might fit in with his.

MICHAEL CASADEVALL: I'll definitely talk with you after. I also have some concerns with the test methodology used by APNIC but before I get into that, I don't want to be hogging all the question time.

FRED BAKER: Paul, how do you want to proceed?

PAUL HOFFMAN: I think that that sounds just fine. And to the last comment about methodologies by APNIC, one is that I want to be clear that if we in this work party do something different than APNIC, that's not necessarily a repudiation of how they're doing it. Basically, when you're trying to get things on the Internet to reply in a certain way so that we can measure them, you're always going to be making different tradeoffs, and I certainly don't want to make it sound it all like that we think that APNIC's choice of tradeoffs is bad. Having said that, one of the things that RSSAC asked us to do is to try to come up with measurements that might be more reproducible and we can even look at how reproducible a methodology that we come up with is, and then APNIC might look at their own reproducibility, for example. So, this might be an interesting academic question.

FRED BAKER: Paul, let me fill in the blank with the thing you didn't say and that is that Geoff Huston is part of the caucus and we've invited him to be part of this. Frankly, we haven't gotten the set of answers that we'd like to have. What Paul is talking about as much as anything else is doing something that we can actually see.

PAUL HOFFMAN: Geoff and João have contributed but have said what they are doing but we would need to be taking some more action on our own before we could say here's how we could compare, for example.

MICHAEL CASADEVALL: Well, my largest concern over JavaScript-based approach is you're running into an issue that both Chrome and Mac OS 10 as whole runs into known as socket racing. Essentially, you've to get the best behavior, the browser – both Chrome and Mac OS 10 as a whole system tries to do both IPv4 and IPv6 connections at the same time and whatever connects first works. That means you get false positive and false negatives both based off the browser and the operating system that user is working, which can give you inconsistent results on how the recursive resolver is working especially because once you decides once methodology is better, it will stop asking for quad A or regular A records if it

knows it's available. That's going to skew some of your dataset. Furthermore, it doesn't give you good information on anything but a quad A or C name because those are essentially the three types of records you can query directly or more or less directly with JavaScript and you can run into unusual problems if MX or TXT records are being unusually manipulated which can happen in certain enterprise, split horizon, or places where DNS censorship is taking place. Things I've encountered in the wild.

PAUL HOFFMAN: I think cataloging those from this work party would be a valuable outcome even if we do not go ahead and create our own system. I know Geoff has talked about some of those in his earlier talks on his test bed, but I think that a reasonable outcome for this work party even if we don't setup a parallel test bed is to list things like that, to list known problems and effects, especially ones that would cause both over and under accounting. So, I'd be happy to talk with you about that and that could be a separate document but it certainly would be a valuable contribution to any future testing that people would be doing where those results then might be picked up by the root server operators.

WES HARDAKER: This is Wes. To be fair, Geoff knows a lot of the limitations of the system and he would certainly love to hear others that you might or anybody might be aware of. Unfortunately, I don't think he's published that list. Paul's right that he has talked about it in previous mechanisms. And specifically, he very much also is aware that it's a very Web-centric mechanism and that there's an awful lot of other use of DNS out there that's not Web-centric, so it's very skewed toward that end. When he looked at the DNSSEC key rolling issues, he was aware that he his overlap with a lot of the other resolvers out there was actually very small and he even disclosed what percentage that was. But I think moving on – I guess I should turn back to you, Paul. Where do you want to go from here?

PAUL HOFFMAN: I also want to hear other people in the room. I don't know if there are other hands being raised and such like that. Let's at least get everyone who has a feeling for, is this useful? What kind of test would be useful? Do we want to put effort behind it because it will take a fair amount of effort to get a delivery mechanism together? It has taken Geoff years of dealing with Google and having things changed just for – I shouldn't keep saying "Geoff." APNIC because at this point it's both Geoff and Joao. Do we want to take the effort to pursue this? And if not, what valuable information can we as work party give back to

RSSAC? Anyways, even if we are not building out a test bed, those are sort of kinds of questions that I would love to hear from folks today.

WES HARDAKER: Those are all fair questions. Since nobody is raising their hand, which I can see the room, I'll go on. There's probably other test metrics that we could get involved including Nick Weaver's – yes, I am blanking on the acronym of his organization that's primarily gamer-focused mechanism, but it still yet another database that he would be willing to probably expose queries from.

One of the things I would like to see out of this work is once the resolver test bed is in place, we can use that as a form of being able to gather fingerprints to how things are being queried and then see if we can take those fingerprints and match them against real traffic or, more specifically, find stuff that is nowhere near fingerprint-wise, that's clear doing something very different. I think one of the biggest pieces of missing information we have about the network at large is, what percentage of resolvers out there are functionally doing DNS resolution in ways that we, more specifically, didn't design for. We know that there's a lot of servers that don't do any caching or very little caching and things like that.

We were talking earlier today in the metrics party of if we're going to design metrics for a system, it's how we believe it should be behaving with respect to caching and other stuff. What percentage of the world at large actually doesn't meet how we've designed this system for? How many other people are misusing it – it's not the right word – but using it in ways that we didn't expect? I don't know how much of a bite we can take on that target, but I think any improvement would be worthwhile.

PAUL HOFFMAN: Okay. And to add, before you go to other folks in the room, on what Wes just said. I should've mentioned that in the list of interesting things or unfortunately interesting things that we know some resolvers do which is caching. Some of that actually can viewed outside of any of this work here by those of us who have access to root server data both DILO data, a day in the life of – I'm sorry, I always use that acronym without spelling it out – this collected, ensured through the [DNSOR] but also – even though I'm not part of the L-Root team, my team in ICANN has access to L-Root traffic data, and many of the root server operators sitting in the room right there do as well.

An example of this is later this year I'm planning on publishing a paper which I'm tentatively calling "Conversations with Root Server Operators," where I look in the data capture and select by

IP address over a course of a few days to also see in the wild with just out with what we have, which of course somebody might be doing things with multiple operators to see what we do. I bring this up specifically because as you look at the most active talkers to L-Root just visually – I haven't done a full reproducible analysis – it's really clear that some of them are not caching at all. So, that would also be interesting for us to be able to get a handle on.

Wes, you were asking what do I want to do next. I want to see if there's enough [oomph] to do some of these, and if not, is there at least [oomph] to be documenting some of these things. And if nothing else, maybe even just document the list of if we had the [oomph], what would we have wanted to measure?

MICHAEL CASADEVALL: I think another important statistic to gather here is, where possible, try and collect information on recursive resolver versions. The two of the more popular packages, BIND and DNSMasq, both respond to the Chaosnet version query. There's not that many recursive resolver packages out there in the wild, so I'm more curious if the caching that is happening is because vendors are shipping things misconfigured or if we are dealing with fundamentally broken DNS software. Because I have a feeling it's more people shipping something like DNSMasq with

an [IDEA] config file or there's a bad default somewhere and we may be able to reduce the amount of traffic to the root zones if we can find and fix this particular default.

FRED BAKER: So, when you say vendors shipping broken code or misconfigured code, you're talking about somebody downloading an open source package, misconfiguring it and burning that to disk or whatever and sending it out, correct?

MICHAEL CASADEVALL: Yes and no. For example, in the case of Ubuntu, we ship BIND and we do not update the major version even secure updates. When we got KSK 2017, we released a stable release update that simply add that to BIND keys. Now, in a lot of home and small enterprise networks, you get like an off-the-shelf router from Verizon or like Asus and so forth, that will handle DNS through your network and can often do basic network management like do dynamic DNS. If those are the types of devices, I'm expecting to be doing horrible things to the root zone and it's either because these vendors are shipping firmware images with bad configuration files or software package that has bad config. I'm more or less speaking from experience because I've built more than a few [embedded] devices that do pretty much what I just described.

PAUL HOFFMAN: I think an interesting thing that we might want to look at – it's sort of a cross between the two tasks that RSSAC asked us, namely look at source code which we're doing in the other test bed and look at resolvers in the wild, is to try to determine – and this may be just wild guessing, but we might be able to elicit some support in the community for a better design – try to determine how many of these resolvers that are being seen by the root servers are actually for what we could categorize as homes. That is, if a home router is configured to be a full recursive resolver or to simply forward or act as a stub sending to a DHCP configured upstream. That may be a valuable piece of information as well especially if we can see trends over time.

MICHAEL CASADEVALL: Given the fact that we have the logs of the root servers, won't it be relatively straight forward to do a reverse DNS lookup of the IP addresses and determine what early chunk of these data it is? In many cases, at least in the United States, residential and small business come out of their own IP pools and can be determined by looking at the pointer records in reverse. I've used this methodology in other DNS-related projects. I don't know if that's viable, if the data is available, but it would be a pretty good way of limiting the scope of knowing where we're

looking from. Is it the small number of devices that are doing something very stupid or is it a lot of devices that are doing something stupid now in it again and just pounding the servers into oblivion?

PAUL HOFFMAN: I'm not in the room so I don't know how many people rolled their eyes when you said, "Wouldn't it be relatively straightforward," but we are talking about the DNS here.

MICHAEL CASADEVALL: I did say relatively, I didn't say straightforward.

PAUL HOFFMAN: Right.

WES HARDAKER: From past studies of looking at the day in the life of the Internet data and trying to do things like that from a worldwide perspective, the data is a lot less clean than is to be desired. I think it could be the safest way to put that. You can certainly try and extract a significant portion of data you believe is clean and trustable, and so what I was alluding to before was when people try to do that in the past they get a much smaller percentage of good data than they were hoping for.

PAUL HOFFMAN: But again, I think that even if we don't move forward, it would be good to catalog things like this and desires so that in the future somebody else even not in the work party wants to set up something like this, we can have set out at least what the root server operators would like to know and some ideas of how this might be designed.

Fred, are there other folks with hands raised?

FRED BAKER: Not that I see, Paul. Well, Michael has raised his hand again.

PAUL HOFFMAN: Okay.

MICHAEL CASADEVALL: I'm going to directly follow up on myself. Sorry if I sound like a broken record. What I'm suggesting is a lot simpler than that. As we essentially take the entire cross-section of IPs in a given day that hit the root servers and then figure out what blocks they're coming from. Because if we can build it out like that, we can see what's coming from mostly residential versus mostly business. Now, I know people are probably are rolling their eyes at this but you can make a broad assumption that most residential IP

blocks should not be doing full recursive resolves against the root server unless there's something wrong with the router. They should be going to their ISP who may or may not be doing something stupid, and we can also then take a cross-section of those IPs and see Port 53 is [world] readable which will give us an idea if it's a public recursive resolver or not. That by looking at the data in that sort of way, we can generate broad terms and hopefully prevent too many false negatives or false positives. If this has been tried before, feel free to ignore me. I'm new here.

PAUL HOFFMANN: You should talk to Wes because of the people in the room, Wes is probably the most research-y who has either done this himself or talked to others. But again, I think we should capture these kinds of things. And maybe even as a way of saying, this is what we want but we don't think that these mechanisms will work but maybe others will or convince us otherwise.

I think to summarize then, if there are no other hands in the room, what I will do is I will make a more complete list of some of these comments and such like that and bring them to the work party as how will this fit into the final report that we are expected to be generating. Then one thing I would like to do is then try to take whatever we do out to the larger research community as a way of even if we didn't do the research

ourselves, of sparring others to say, "Oh, they should've done this. Well, I have ample free time, I'll do it," or things like that.

By the way, thank for the contributions from all of you. This was more discussion than I expected but I think that that's all I have for the work party. Fred, if you don't have anything more, maybe we should turn this back over to the metrics folks who still have some more open questions.

FRED BAKER: We'll do that in a moment but Russ wants to get in and Brad wants to get in.

RUSS MUNDY: Hi, Paul. Yeah. This is Russ. One of the things that comes to mind in terms of both the history of the work party and the work itself and recent events that I don't know if there's any way we can get a strong correlation that is thinking about this and what occurred back in January when the 2010 revoked bit was set and we started to see these odd traffic anomalies that were almost certainly coming from resolvers. I don't know if there's any way to identify a need to characterize events that have occurred that we don't necessarily have a good explanation for is coming from the resolver world. That's one of the recent ones that come to mind. And I'm not sure who all did the looking and digging into

that. Did we ever come up with a good answer or reasonable answer for what was the cause of that spiking?

DUANE WESSELS: Yes. Wes did a lot of work on this. He found particular versions of bind that behave badly when they get a revoked bit and wrote up quite extensively about it. I think in that case, we do have a pretty good explanation.

RUSS MUNDY: So it was a unique characteristic that was directly software-related that isn't something that would kind of fit into the bigger boxes of what do different resolvers do in broad case instances. So it was a very point case kind of thing.

DUANE WESSELS: Yeah.

RUSS MUNDY: Okay. Thanks.

WES HARDAKER: Well, to clarify further, I identified a few issues with a few software components over the course of probably a year, some of which have not been fixed, some of which have not been fully

studied to figure out what versions they exist in. I believe there's at least two bugs in at least one software package, not one, and we haven't caught them all. So, in five years we may see it again if nobody cleans up that particular issue or is able to reproduce it under a lab condition.

That being said, at no time did I ever fully attribute every measured packet that we saw down to that particular problem. So though I'm confident I found a couple of problems that were definitely contributing, I don't even know the percentage of which they were contributing. The VPN software, I could measure the percentage because we saw the change, but the other one, because no fix was deployed by the time that the traffic levels changed, we really don't have a clue.

RUSS MUNDY: What I was really having in mind in terms of the work for the output from this work party is if we could perhaps include at least a listing of unexpected behaviors that have been seen recently that some could be attributed, some couldn't be, but did in fact result in a significant impact on the root servers. Is that a reasonable thing to think about including in this, Paul?

PAUL HOFFMAN: It's reasonable. I don't know of any instance other than the one you brought up. As Wes and Duane were indentifying … We can maybe say there seems to be something happening to the root servers but identifying where it's coming from is quite difficult because we're not talking about versions. We're talking about versions and specific configurations. Throughout the KSK rollover when we were finding unfortunate things, it was almost always a combination of a version and a particular configuration, some of which were the recommended configurations and some of them were configurations that the vendors would say, "No one in the right mind would ever do that. Oh my gosh, look, a lot of people did it." So, maybe not as part of the work party – this work party looking for anomalies like that, but it seems like something that RSSAC might ask a different party to do is to try to find groups of anomalous behavior and hopefully the outputs of this work party might help figure out why.

FRED BAKER: And now to that end, you have a GitHub repository. Do you consider that more or less complete?

PAUL HOFFMAN: No. As I said, at the last meeting, it's definitely not. I still have some work to do. Paul Muchene who is in the room there has

been opening issues and such like that. No, there's still a few open issues on that. Plus I haven't actually run the tests that will get interesting output. I'm expecting to do that in the next couple of months.

FRED BAKER: Okay. It seems like one of those things that having built some software, it would be nice to use it at some point and put the results in that paper when we produce it.

Okay, do we have anything more for Paul at this point? Mike?

MICHAEL CASADEVALL: Just as a final follow up to that, if we have identified versions of bind or other DNS software, please contact me because I can help get those fixes backported into most major Linux distributions and fixed as stable released updates. Also for purposes of testing, I have a series of Docker images that emulate the entire root zone with its own DNSSEC key which may be useful for reproducing certain test cases. That's available in GetHub and I can get that to anyone who's interested after the meeting.

FRED BAKER: Okay. With that, I think we're done as far as the resolver meeting is going. [Ozan], maybe what you want to do then is turn off this recording and turn it back on for the metrics people.

PAUL HOFFMAN: Actually, wait. Before you do that, if you turn off a recording and turn it on, all of us remote people are going to get kicked off and I don't think we know which link to hit to come back in.

FRED BAKER: Okay. I'm just trying to not confuse the metrics people. Is there an easy way to not confuse the metrics people?

OZAN SAHIN: It's going to be the same link that you connected for this meeting.

PAUL HOFFMAN: Okay.

FRED BAKER: So, in asking for that to be shut down, am I making things more difficult for you?

UNIDENTIFIED MALE: I just need to make sure that I can divide and put the recordings properly together in that case [inaudible].

FRED BAKER: In that case, let's not do that. Let's not kick people off. Duane and Russ, your party.

UNIDENTIFIED MALE: I can make a note in the records you're saying like at this point you're [inaudible] previous session so people will have that [inaudible].

DUANE WESSELS: Alright, we're ready to start? Okay. I'll ask Ozan to put the Metrics Work Party document back in the Zoom room and if you want to follow along, you should have the URL if you're in the caucus.

When we ended before the last session, we had made it to the end of Section 3. Section 3 is about general applicable aspects of the metrics. Section 4 gets into some specific measurements and metrics for root servers.

The first one talks about root server availability and this is broken down by v4-v6 TCP and UDP. One thing that's maybe a little different since some people have heard about the metrics

is – for the most part, now we're proposing that all of the measurements be done at five-minute intervals. Previously there was sort of a mixture. There were some at I think one minute, some at five minutes. There is still one of them I think the freshness or staleness. We talk about doing it every hour just due to the way that one works. But all the rest are at five-minute intervals and did that for simplicity to keep them all consistent. So that's reflected here.

Can you scroll down a little bit? There's a section that talks about – yeah, Wes has a comment. Wes has left the room so I'll have to read his comment. But the section talks about – it says the measurements have a timeout of four seconds and for response received within the timeout value, the root server is considered to have been available. Wes's comment is: "What if that response indicates some kind of error?"

Off the top of my head, my opinion would be to not address that case specifically at this time, that we would still count that as availability, but maybe for future work, we would delve more into that if we need to separate that out. I don't know. I'm open to suggestions. I guess my primary motivation is to just keep things simple and get something out the door rather than try to solve all the problems on the first go around.

RUSS MUNDY: Well, I guess Wes has stepped away for a moment. But in a way, I guess this raises another question. When you're collecting data for a particular metric and something occurs that is incorrect or problematic from another metric's perspective, do we want to try to tackle that at all?

I'm in agreement with you, Duane, that we try to keep this first cut through as straightforward and simple as we can, and if it's something that's talking about availability and a response, if you get a response, you get a response, and we go forward from there.

DUANE WESSELS: Yeah, I agree. For the most part, these metrics are designed such that the availability – one that's sort of the easiest one to do. That's the first bar. The remaining metrics which say, for example, you don't include timeouts in the remaining ones. So if you're worried about timeouts, you look at availability. You ignore timeouts and those sort of errors for some of the other things. But I haven't previewed these comments before so I'm just reading them the same time I'm reading them to you. I haven't had a lot of time to think about it. But that's my initial reaction.

RUSS MUNDY: It does seem to me that that's the approach that we need to take, and if people in the implementation phase determine that it's more efficient to combine the data that you collect that was primarily for one metric and use that data for another metric then if there was some sort of problem or anomaly for the other metric, it gets noted in that. But from just a pure metrics perspective, keep them separate in terms of the description.

DUANE WESSELS: Yes.

RUSS MUNDY: Does anybody see a problem with that?

DUANE WESSELS: Okay. Another thing I want to call out specifically is in the section that's a little bit lower, it talks about aggregation. Again, this is something that I think is now consistent for all of the metrics. They all specify an aggregation period of one day.

Previously there was something here that talks about – you could aggregate this over any time period but you had to have some kind of minimum number of measurements to make it a valid aggregation. But I felt again that was getting complex and confusing. Now it says when you go to reports, for example,

availability, you always report it over a one day period. I think elsewhere up in the document in the general section it talks about how timestamps are all based on UTC, so this would be a UTC 24-hour period, not a local time zone period, that kind of thing.

Let's scroll down some more to the example. This is what I was getting at before. This table kind of matches the style of the IANA performance reports. Those are actually done monthly so their numbers are reported on a monthly basis, not on a daily basis such as we might be doing here but this is how it looks. There's a heading that says this is X-Root IPv4 UDP Availability for a certain date. There's a line that repeats availability and it says here 99.7% which is in green, meaning that it has met the threshold. Now, please keep in mind these are not actual proposed thresholds. These are just examples to illustrate the format of the table. These are not intended to be actual thresholds that would be used. But then the line below says the threshold minimum value as proposed here by Wes was 99%, so since it's below that, it's green. Then it includes the count of how many measurements were included in that data. Ozan? Remote?

OZAN SAHIN: Paul has his hand up in the Zoon room. Paul?

PAUL HOFFMAN: Duane, you just said something that I think is not going to be true in the future, which is what we're going to aggregate by one day. But people looking at this data will then aggregate those aggregates. So I think what we are likely to see is – again, only thinking here about the availability metric – that many people will say X-Root had 100% availability 95% of the days. Really, whether we put a threshold or not, which is still an open question on it. Even if we report these, a simplifying thing would be how many days were they fully available? And I don't think there's anything we can do to prevent people from doing that but I think we should be aware of that. Thanks.

DUANE WESSELS: I agree. We can't prevent people from doing that. I guess the question that we face is, what do we recommend that this data looks like coming out of the measurement system or published by the organization that's responsible with publishing these metrics? You suggested something – I heard you describe it as simple but to me it sounded a little more complicated. You said 100% availability for some number of days. Is that what I heard, Paul?

PAUL HOFFMAN: Yes. If we report – because you said that the aggregation that we're proposing to report here is by day – somebody is going to

re-aggregate those to – or let's just say that we do have a threshold value, somebody is going to say X-Root missed its threshold on N percent of days. We might even do that ourselves. Even though we might be aggregating by the days, the threshold might be you must be this good. And it's okay if you fail once in a month but you can only fail once in a month. So that's a second level of aggregation that will happen. So even though you say here what the aggregation is going to be for the metric, anything which is a real threshold or even a perceived threshold will be an aggregation of aggregations. Does that make sense?

DUANE WESSELS: I suppose. I'm not sure I'm opposed to that or that it's a problem though. Do you see that as a problem?

PAUL HOFFMAN: I'm sorry. I do see there's a problem if we don't have thresholds because people will make up thresholds if we don't. But if we have thresholds then it either matches or doesn't match the thresholds.

DUANE WESSELS: Okay. So is your proposal then just to make sure that we have thresholds or would you propose something different to address the concern?

PAUL HOFFMAN: I'm waiting until we talk about thresholds to have that discussion. I don't know here and I'm not saying that we should make any change until what you have here. In fact, I think what you have here is fine. But for availability, as we discovered in the discussion earlier today, the thresholds are going to be a little bit dicey for some operators who in fact will miss a normal availability window but still be in fact serving the root well within the region that they wanted to be serving it in. So a few people had suggested we should let root server operators pick what kind of thresholds are applied to them if they're trying to be very region-specific and such. So that whole discussion is going to answer the question of, do I think we should have thresholds here or not?

DUANE WESSELS: Okay, yeah. That was an interesting suggestion from earlier today that there could be different SLA thresholds for different people. I'm not sure that that's simply something the work party should be getting into. I don't know. But apparently now is the time because we have hands up in the air.

RUSS MUNDY: Well, we've got a couple of hands. I think I was first here. My reading of our work party charter for this is we're to lay out what the work party believes for what [will become] an RSSAC view of what the threshold value is for these metrics. It might or it might not be the threshold that would be used by some group that results from the community process associated with 37. They might use, they might not. It might be an SOA, it might not. Our job though is to try get I think a singular reasonable threshold for each of these, and then if it's used beyond RSSAC then so be it. But for us, there's not one for every root server operator.

JEFF OSBORN: I'm being concerned that we're measuring something and what gets measured gets done. This really runs counter to much of what we want to do in terms of diversity. So if you're working on a grade and you're close, you're going to put your next six servers in Manhattan when the need may be in Sumatra or [inaudible] or someone else. To the degree that we have a big asterisk here and explain to somebody we're describing this to that please don't use a single score that's better. But this is like buying a baby carriage because it's faster. It's a silly metric when used poorly.

UNIDENTIFIED MALE: I'm concerned with how availability is handled for these faraway hard-to-reach root sites. But I think things as written will work fine here. If I have a site off on a corner that goes down quite a bit, I'm going to have other nearby sites that from global monitor's perspective are still reachable. I'm a little concerned about the availability metric for my local sites that lose power for a day or two at a time. But as written here, it's the amount of time that the – it's not divided by a fixed number of [inaudible] a day, it's the number of measurements that were taken. As long as the measurement device is down at the same time, the root server device is down, the availability won't be held against it. So, local ones work and remote ones seem to work.

DUANE WESSELS: So if I had put in this example that the availability minimum threshold value was 10%, would we be having this discussion? You had to be up 10% of the time. Would we be talking about root servers in faraway places? Is it just because I picked a stupid value here? I mean I feel like we're having the threshold discussion right now rather than how to do a measurement and how to present a metric.

JEFF OSBORN: Since nobody else is saying anything, that's probably a very valid point. You're right.

UNIDENTIFIED MALE: We'll probably have another discussion.

DUANE WESSELS: So I should use maybe imaginary numbers here. It's 99 I.

BRAD VERD: Still it begs the question. When is the threshold discussion going to happen? And I will add that in my opinion, there should not be different numbers for different servers. There should be the same SLA for the service across the board, and that has nothing to do with diversity – nothing.

OZAN SAHIN: Paul Hoffman, please go ahead.

PAUL HOFFMAN: I'm going to bring up again the fact that if what we are caring about as the community is how well do the root servers respond to resolvers and resolvers are going to pick the best root server, any question of thresholds for availability are going to be very, very weird because a root server operator that has 15 maybe mostly unavailable to you as a probe servers and just one that is well connected, when you are testing it, it's going to go to that one. So, when you get this availability number, again, we're

saying we're only looking at the RSO. That RSO is seemingly available to you even though many people who were looking at the instance would say, "No. The vast majority in our instances are unavailable." Again, maybe we postpone this until the threshold discussion but I think we need to keep remembering that resolvers are doing some of the work for us on any of these aggregations.

BRAD VERD: I don't disagree with that statement. However, we're not here to measure the resolvers. We're here to measure the root server system and the individual service provided by each root server operator. So if the users are experiencing something different through the resolvers then that's not indicative of the health of the root server system, that's indicative of the health of the resolver system. I thought that when we talked earlier today – and I'm sorry I missed part of the discussion which was I guess a key part of the discussion, I apologize for that – there was a discussion about whether the probes would use local resolvers or have it built in on their own and I don't remember where that ended, but on their own which would basically be with it, we wouldn't be testing the resolver system so that a probe would go directly to a server. I think we need to be careful that on … again, try not to overengineer this thing, try not to overbuild this

thing, but trying to get an availability number from a set of probes.

LARS-JOHAN LIMAN: I would like to augment that. Try to get two availability numbers because we need to keep an eye on measuring individual server operators and measuring the system as a whole. That's the thing that keeps coming back and biting us.

BRAD VERD: That is precisely what I just said. I said the root server system and the service provided by the root server operators.

LARS-JOHAN LIMAN: I heard that. Then you finished with, "Let's try to find a number." I don't want "a number," I want two numbers.

I fully agree with you, Brad, but on the other point that we need the same numbers for all the service operators … whatever we measure should be the same in all places. We cannot negotiate separate numbers for separate corners of this measurement system because then the system will again be rigged.

DUANE WESSELS: Well, I think to be clear, we can recommend that they not be the same. We can recommend this sort of thing but it's probably not

going to be up to this group to decide that, right? It's going to be this other function eventually.

RUSS MUNDY: This group can decide and if RSSAC endorses it then we can say this is the RSSAC position, but if others outside of RSSAC decide to do something different, we don't have any control on that. But from what RSSAC eventually says, that's what we're providing the input for.

BRAD VERD: I hope this group can make a decision. Quite honestly, what you just said, Duane, I hope we never come to. We need to make a decision.

RUSS MUNDY: I'd like to make a little bit of a response to what Paul said earlier and correctly reflecting how regular resolvers work today and that is many of them work differently and they'll pick the one that is the fastest responding. All of this is good, but part of I think the decision that we made this morning, we need to confirm in the mailing list that we want the resolver functionality to be part of the test point or the probe or whatever we end up calling it and that point won't be doing caching, so each one will go out and make its query.

The part that I don't know is how hard it will be to say – and we can say it in our document – that it should not be doing the kind of preferential decisions in terms of which root server it goes to based on past performance. It should go make the queries in accordance with whatever the algorithm is for making queries, not in accordance for what's normally done in resolvers of what's the fastest. I think that's something we can say. I don't know how hard it would be to implement.

OZAN SAHIN: Paul's hand is up. Paul?

PAUL HOFFMAN: I'm sorry if I was unclear with what I said just a little bit before about resolver and root server selection. I was thinking more locally than globally – and I know we're sort of half talking about thresholds here – but there are easy to imagine instances where a root server operator will get low scores even with aggregation on a one day level for availability, which is exactly the metric we're talking about now, and yet still be contributing to the root server system in their local region.

We are now seeing – I think currently there's at least four countries which are blocking Internet access to the outside world but there is still Internet access within the country.

Because none of the 12 operators of the 13 letters today are doing this, but a root server operator in the future who might be specifically trying to set up a root server system within those countries for the very reason of historically our government has been an idiot and has shut us down would be getting fairly bad rates on availability but would actually be serving the root just fine within the country. When we have the threshold discussion for availability, if that goes against those folks then we won't see as many of them set it up, and to me that would be sad. But again, this is a discussion for thresholds, not for this.

BRAD VERD: Again, I apologize if I missed something earlier this morning. Paul, regarding your comment there, are you subscribing or are you implying that every instance would be monitored for availability?

PAUL HOFFMAN: No. Definitely not. That's what I'm saying, that an operator who set up … and again, according to RSSAC42, I think that a new operator could in fact not even be anycast, which I think is a bad idea. But an operator who is running a single instance or runs a small anycast network within one very small geopolitical region from the outside world could be failing on availability. Well, if you had put the probe inside, which you wouldn't be able to, to

be able to run if there was a government mandated shutdown to the outside world.

BRAD VERD: I kind of feel like this scenario that we're coming up with here is an exception scenario – and I don't believe we should be engineering two exceptions. I'll state this: I believe that the root server is a global service. Therefore, it should be monitored globally – not locally, not geopolitically. It should be monitored globally. So availability should be from a global perspective, not from a geopolitical perspective.

It's interesting to me that when we talk about exceptions about my servers somewhere in Africa or somewhere down in Australia or some far off region, we don't talk about our instances that are within China or somewhere else, which is the geopolitical piece that is being talked about. So it's an interesting conversation. I don't think we should be focusing on the exceptions. I think we should focus on the global service.

LARS-JOHAN LIMAN: Plus one. I just want to reinforce what he said. This is a global service and if a future government system were to look at adding root servers, I hope they would – I wouldn't say disregard – but definitely give substantially lower point to an applicant to only

serve a certain region. That said, there are exceptions and they need to be handled. But I think, as you said, we should not design for that with this system. There will be separate provisioning for root servers and therefore there should also be separate monitoring, which is not part of this discussion.

MICHAEL CASADEVALL: My one concern with this is to the way anycast routing works. Where you are doing the measurement physically is going to affect how the availability scores work. For example, if you're within China, they can easily advertise BGP routes for their own set of L servers if they choose to run them that may or may not reflect on what the general routing network is. So your definition of availability very much also depends on your definition of where you're doing it. You have two variables that constantly change. If you take all the readings from a single location, you'll get one number of availability. But if you take it from multiple locations from different points in the Internet, due to the way BGP routing works and the route to the fastest end point, you could skew your results in very ugly ways. So I don't think the issue is moot but I think it needs to be expanded that the BGP route needs to be taken into account from the source address. Fortunately, you can get that information from looking so it could be automated.

DUANE WESSELS: Did you still want to make a comment, Brad?

BRAD VERD: This is going to add – if we take the geopolitical stuff aside, the example I use is that I have instances that are local only, meaning they are serving a closed network or a closed society, much like a geopolitical area but they're not, and I don't expect them to be included in any availability metric. So I might be down on the availability side but still be serving them in this local world but I'm not going to be arguing that my service is up. Do you see what I'm saying?

DUANE WESSELS: Alright, we've got five minutes left. Go ahead, Russ.

RUSS MUNDY: Let me try to summarize this discussion. How do we describe this and what are we specifically trying to describe and get included in the document here? Does anyone disagree that in fact what we want to have to find in the document is a single percentage value for each of the lines that are in this table that will be applicable to all of the RSO individual operations as we've talked about it before? If there's variations from that, they

get made elsewhere. But from an RSSAC perspective, this is what we're striving for. It is to get this table with a single set of values for each of these. Is there an agreement with that? No disagreement. Okay, good.

FRED BAKER: Speaking strictly for myself, yeah, I agree with that. That's fine. We will get into interesting things with the servers that are in, no export zones, which I think is what you were talking about.

RUSS MUNDY: Right. The exceptions to that need to be handled separately and outside of what this document says, whatever the appropriate process is.

BRAD VERD: I don't know if they need to be handled. That's what I'm saying. So if I have private peering and I provide that service privately, that's not necessarily included in my availability metric for the global service.

RUSS MUNDY: What I meant by "handled" was it could include ignored, not part of it at all.

BRAD VERD: Right.

DUANE WESSELS: Alright, I'm not sure what to do with the remaining three minutes. I don't really feel like timing in to the next section. Maybe we should just stop and claim semi victory. I don't know.

RUSS MUNDY: I think we've got progress in terms of reaching good agreement on what we're trying to do. The material we have here I think is good, it needs tuned up but I think everybody is okay with the concepts that are in the current document. Am I being overly optimistic there?

BRAD VERD: Optimistic is a great way to approach it. I obviously didn't hear the follow-up to my grenade that I dropped in the room this morning when I left but I'm sure I'll hear about it from others.

DUANE WESSELS: One thing that the work party could use some input on I think is the staleness/freshness. In Section 4.4 there are three proposed or two and a half proposed methods for measuring this. This is sort of a hard one to measure. So if people would like to take a

look at that and maybe voice an opinion for one over the other, I think that would be very helpful.

` This is a little bit of a difficult one to pull off because you need somebody. The probes or some central system needs to know the ground truth to compare to or you have to do some kind of consensus on what is everyone else serving in terms of up-to-date content. So, this section is a little bit longer than some of the other ones that talks in various ways about how you can collect the data and also deal with the situation where the serial number changes in the middle of your measurements and handling that. So I would appreciate if people take a look at that.

Russ was suggesting that we can talk more about that after the RSSAC formal meeting if we want to. If we do that though, I suppose we would need to invite the caucus back in somehow.

BRAD VERD: When is the next work party meeting?

DUANE WESSELS: I don't believe it is scheduled to this time but we've been meeting on basically every two weeks schedule. So we'll try to keep that up. I'll work with Steve and Ozan to schedule something. Usually Paul goes out and we schedule it.

FRED BAKER: Two weeks from now would be July 9, am I correct on that? The week containing July 9. Then the following one, I could imagine being at the IETF meeting. Is that [inaudible] that makes sense?

DUANE WESSELS: Yeah, that sounds about right. Yeah. That's the schedule we've been trying to keep.

BRAD VERD: If I may, I just want to thank the Chairs and everybody who has contributed. It is some really good content and I know this is a tough conversation. So, thank you.

DUANE WESSELS: Just to be clear, the work party would like a meeting for the week of July 9. Is that when staff should schedule one? Okay, thanks.

FRED BAKER: You're running this, so yeah. Okay, fine.

RUSS MUNDY: In that case, adjourned.

FRED BAKER: We're coming back for the RSSAC meeting in 14 minutes. So, run to the bathroom. Do what you're going to do.

[END OF TRANSCRIPTION]