
UNIDENTIFIED MALE: [inaudible] and it will be recorded. So thank you. Over to you, Duane and Russ.

DUANE WESSELS: Okay, thanks. So welcome, everyone. Welcome, work party members. Steve just said this will be treated as an extended work party meeting. If you've looked at the agenda or had a chance to look at the agenda, we have a lot of time now on the agenda to discuss metrics. I believe it's the next nine sessions or so. So almost the rest of today and then all of tomorrow and a little bit on Thursday.

Yeah, welcome. So we're going to start off with bringing everyone up to date on what the work party has been up to so far. Then, in the rest of the time in this slot, we hope to have a discussion around the question which is how does the work of the Metrics Party tie in or relate to RSSAC037 because that's probably going to drive the [future] discussion in the work party.

So Russ and I have agreed here that I'll start off with what's been going on to date. If my memory and calendar is correct, the work party has met three times now. The first meeting was very introductory, going over the statement of work and that sort of thing. That's where we also – Russ and I – volunteered to co-chair this work party, and we did that sort of administrivia.

The second meeting was quite a bit more productive. We had asked work party members to fill out this spreadsheet and provide their

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.

interpretations, if you will, of what the statement of work was asking. We got some very good results from that, and I think it helped a lot to drive the discussion. If nothing else, it showed places where people had very different interpretations of what was being asked. So we could spend time on those things.

Then we had a third meeting after that, which we covered this in more detail and discussed the outstanding issues amongst ourselves. From that, it led to this document that's up on the screen. We're not going to really talk about this document probably right now, but you're looking at this Google document over the next few sessions, where we're going to present where we're at so far with the work party.

Another thing that has happened within the work party is that, in addition to the feedback from the work party members, requests went out to the root server operators to say what sort of things are the operators currently alerting on or monitoring for themselves. We got some good responses from that. Those responses are actually summarized in this document at the end. We don't need to look at them right now, but if you can look at them or on your own – or maybe we'll get to it later – those are in this document.

Then the staff also did a survey of the existing RSSAC documents that are relevant, like 001, 002 – I think 024 is the expectations one – and also the RFCs. They pulled out things that may be relevant for the work party to discuss that can be measured protocol-wise or expectation-wise.

I don't believe that is represented in this document but we do have it separately. If anyone wants to see that, we can point you to that.

All right? Everyone good so far?

So a lot of what the work party has been focused on is – we tried to divide into three broad areas. One is metrics relating to individual operators. Next is metrics relating to what until now I would call the system as a whole. I'm going to try to stop using the phrase "system as a whole" because I think we may be getting a little bit stuck on the word "whole." But I'm going to try to say "to the root server system." So we've got operators on one hand. We've got the system on another hand. And then the third part is something from the statement of work, where it asked to clarify bandwidth [packets] and queries per second. We're going to actually devote the entire session after the next to this BPQ topic.

Anything to add to that, Russ? Are we good?

Good. So the BPQ is actually – one of the reasons it's going to be next is because it's something that the work party has struggled with a little bit. We need to talk through the extent to which the work party should be working on this for reasons that it's a little bit different than some of the other things.

Another thing we've been a little bit stuck on in the work party is – well, as I mentioned before – how do the metrics tie into the 037 document. That's one thing that we need to spend time on today. Another one is we have some, I guess, ambiguity or uncertainty about what it means for the root server system to be X. What does it mean for the root

server to be online? What does it mean for the root server system to be correct? So we need to get some agreement and clarity on that, I think, before the work party can finish some of those metrics.

In both the metrics, where we're talking about the RSOs and the RSS, we've settled on four particular metrics at this point to look at. One is availability. Two is latency. Three is staleness, and four is correctness. So availability is pretty straightforward. I think, in the statement of the work, [it maybe used] the word "online," but it's essentially availability. The latency and staleness in the statement of work came out of a bullet point that talked about being timely, but we split that up into two separate things: one being query response latency and two being staleness of data being served, like SOA serial numbers, if you will.

RUSS MUNDY:

Just a quick comment. On the availability, it's probably the least contentious one, but we've had a wide range of input from work party members as to how you should go about assessing it. So it will probably take maybe not as much as discussion as the rest to get to an end point, but it's not discussion-free.

DUANE WESSELS:

Yeah. So the fourth one was correctness. In the work party's discussions so far, this has really focused on DNSSEC validation. So that is sort of the way of determining that a response or that some behavior is correct.

We also have some of what I would call more protocol metrics; making sure that UDP works, making sure that TCP works, making sure that

check sums are on. That kind of stuff. We have those in the document but, to be honest, the work party has not spent a lot of time talking about those at this point.

Can I take any questions or can we take any questions at this point about what the work party has been up to so far? After that, I think we can move on to addressing some of the particular issues.

UNIDENTIFIED MALE: [inaudible] statement of work? I'm trying to remember back when this started what the end result was to be. From our discussions this morning, I was thinking it should be pretty much a descriptor of where the bar is and what the SLAs are going to measure.

So are those four items likely to be the totality of what gets tested?

DUANE WESSELS: This gets to the heart of the first question that I think we need to talk about – the extent to which the Metrics Work Party relates to RSSAC037 – because I see it as not a one-to-one mapping. So, for example, the statement of work does not say, “Go out and define these metrics that can be used to judge operators in the new governance model.” It’s more broad than that, like we might be doing this work even if we never had RSSAC037 to refer to. This is work that has needed to be done for a long time, so [we] might be defining metrics that—

UNIDENTIFIED MALE: It could be used by 037.

DUANE WESSELS: It could be used by 037—

UNIDENTIFIED MALE: [inaudible] and easily be used by 037.

DUANE WESSELS: Right.

UNIDENTIFIED MALE: That was the whole point of maybe starting this work now and not wait for 037 to be implemented: if we came up with the technical accountability piece, it should just plug right into 037 and provide you with what I think you're asking for.

DUANE WESSELS: Well, it just seems like this should be at least a superset. And if it's those four items, what I'm asking is, are there more things we're going to wish were looked at when it comes time to look for a bar?

UNIDENTIFIED MALE: I can't answer that. Yeah, I—

DUANE WESSELS: Yeah. More thing I don't know. I think the work party is a little bit cognizant of biting off just the right amount this time. Maybe we'll add

things later if we need to, but let's get something done rather than boiling the ocean right away.

The statement of work, at least in my recollection of it, doesn't talk about SLEs, SLAs. It doesn't use that phrasing. Maybe Steve or Andrew – do we have the statement of work that we can put up on screen?

WES HARDAKER:

I was going to summarize the four things that we put into the scope for it. I pulled it up really quickly. Basically, it says that we should do four things. We should define externally visible metrics that can be measured, define measurements for the minimum level of service – so there's your bar – refine what BPQ is and spell that out better, and then update documents accordingly, which includes 001 as well as RFC 7720, potentially.

UNIDENTIFIED MALE:

That pretty much covers what I think you're looking for.

DUANE WESSELS:

I think so.

WES HARDAKER:

That's the longer version.

DUANE WESSELS: “Define the measurements to ensure that” ... right. So maybe this is a fine point, but it says, “Define the measurements to ensure that we’re meeting a minimum level performance.” I think the point that we want to get to is we want to say where the bar should be in addition to what the measurements are. How do you measure the bar and where the bar should be. We want to do both of those things. So, as we go through these sessions really tomorrow, you’ll see that a lot of TBD (To Be Decideds) are where do we set this bar? Assuming we have agreement on how to measure the bar, then where do we set it? That’s what a lot of the discussion will be tomorrow.

[FRED BAKER]: May I make a nitpicky point? It says, “Define measurements to ensure ... When you read that, you said, “Define the measurements,” and “the” actually turns out to be an important word because, if you’re defining “the” measurements, then you have enough of the [“Man, they’re] the right ones.” I think what we’re asking the work party to do is define a set of measurements [inaudible] might be a good first guess. That might change later.

DUANE WESSELS: Yes, I very much agree with that.

RUSS MUNDY: And one of the other pieces that has been discussed fairly extensively in the approach to, if you will, achieving those measurements. Is it better to try to look at the measurements that are either being done today or

desired to be done on individual RSO activities and combine that up to a [system] view, or is it better to try to determine what the right system description of those measurements are and then spread them out to the individual RSOs?

BRAD VERD:

I'll just add a comment because [inaudible] further say a couple times, since you guys have started here, this concern about the system as a whole, I guess. I guess, if I think about that, I can see where you could begin to rathole. I guess the intent – I'm looking at Wes right now – there was maybe it needs to be rewritten or changed to say "the service as a whole," and not the – because – and I'm speculating because I've not been on the call when this happened, but I was on a couple of them but not all of them. I'm speculating, if somebody says "the service as a whole," that means every device that makes up the service needs to be monitored or needs to be reported, and that's not necessarily what I think the goal was when you talk about the system as a whole.

The system as a whole, which I think we were trying to get at, was the black box that is the root service.

DUANE WESSELS:

Yeah.

WES HARDAKER:

I agree.

DUANE WESSELS: And in the document that we'll be going through tomorrow mostly, I think the "as a whole" has been removed from that. So it really does try to talk about "the service."

BRAD VERD: Okay.

DUANE WESSELS: But that's certainly the direction I want to go as well.

RUSS MUNDY: Yeah. And I think we're in alignment, Brad, as far as "the service," versus every individual tiny and little piece that makes up "the system" as a whole. But the discussion challenge still exists in terms of trying to figure out what is the appropriate way to describe the service for the metrics.

BRAD VERD: That's a religious question that you're never going to answer.

RUSS MUNDY: Well, in a way, it would be very much like when another zone, whether it's a .com, or [tislabs.com], or versign.com, looks at their name service for their zone. They want to have a service of a particular type, particular quality, for the zone itself. So at least my mental picture of this is the equivalent of that for the root zone, even though the actual service provisioning is done by different organizations.

BRAD VERD: I guess what I'm fearful and what I'm referring to by saying it's a religious question is how you do that monitoring because there are so many different ways to do it, different ways to implement it, different tools to use, that that could become a really scary rathole of time and effort that you never come out of.

RUSS MUNDY: Yeah. The measurement system is part of what's being discussed.

DUANE WESSELS: So I wanted to give a chance to anyone on the phone to chime in on this discussion right now.

BRAD VERD: [inaudible]. It wasn't me.

STEVE SHENG: Are there any inputs from the people from the bridge? This is Steve.

No.

DUANE WESSELS: Okay. So, on the agenda, we have a lot of time in this slot to talk about how these metrics relate to 037. I don't know that we've [sufficiently] done that. We talked a little bit about it, Jeff, with your question, but I guess I would reiterate my point that, for example, it would be okay for

this work party to define or discuss metrics that maybe weren't [stipulated] by 037. Maybe they're something we thought of afterwards, or they don't necessarily have to fit in there, and maybe vice-versa; the work party doesn't have to address everything that's also mentioned in 037. So I see them a related but not one-to-one mapping, if that makes sense.

STEVE SHENG: Microphone, please.

UNIDENTIFIED MALE: I'm sorry, why don't you see it as one-to-one mapping? If you were talking about a superset I can understand that, but if you're talking about a Venn diagram but only a couple parts that cross, then why are we doing this?

DUANE WESSELS: Which one is the superset of which one in your mind?

UNIDENTIFIED MALE: This should be a superset of whatever we're going to end up using to test whether somebody's tall enough or what "good" looks like. So what good looks like should be, at worst, a subset of this, I would think, because if we're missing something, why not add it now?

So this seems like an interesting time right now to look at the scope here and say – or maybe this is a rathole – what's a test we could

imagine? We could propose a test. To pass, you'd have to have 10,000 packets per [fortnight] with a latency of less than a month. Just start with something because that's what we're going to end up doing: applying this. Right?

DUANE WESSELS:

Yeah. So I agree with your answer that the work party metrics should be a superset of the 037 rather than the other way around. I think, as long as we – everything we can think of that's not impossibly hard I think the work party should be willing to take on. So to the extent that there are things that we may know about now that aren't written down, we should include them or we should write them down and get them done.

So the question here becomes, is this sufficient?

Yeah, and maybe we won't know until we're closer to the end of the work party's output. At this point, it's still a little bit early, I would say, but definitely we should keep that in mind.

RUSS MUNDY:

I think where the work party is at this point is, as far as what there's been agreement in the work party, that the primary areas that Duane mentioned earlier are the ones that will be at least the minimal set of what the metrics are. So whether that's done on an individual RSO basis or some aggregated basis in some manner, we haven't even gotten to that question yet, but the basic metrics themselves.

I agree, Jeff. This is an excellent time to discuss whether or not what we're looking at this point is sufficient.

So let's have thoughts and inputs from both those on the call and those in the room as to the sufficiency of the basic set that we're ...

FRED BAKER: Remind me what that basic set is.

DUANE WESSELS: So, at this point, we've been focusing a lot on four things that are measurable: availability, latency, staleness, and correctness. We sort of have metrics for all of those, both as individual operators and for the service.

FRED BAKER: Okay. I have a question on latency. I agree that it's something that can be measured, but does that make it the right measurement? For example, suppose I was standing in Beijing and the service I was looking at at the moment happened to be B. There's a good chance I would have a very high latency. It would have nothing to do with B. It'd have everything to do with that I was in Beijing.

What aspect of latency is an appropriate measure in this context?

DUANE WESSELS: A couple thoughts on that. Now might be a good time to talk a little bit about this measurement platform, which is under this stipulation section. We had this originally on the agenda in its own slot but it got

pushed to the end, so I'll just cover it briefly. If we want to take a lot more about it, we certainly can.

FRED BAKER: [inaudible]

DUANE WESSELS: Yeah. In the work party's discussions, we've always been working with the assumption that there will be a yet-to-be-defined platform from which measurements can be made. And they would be distributed. They wouldn't all be in the same place. They would be distributed around the Internet with some amount of distributedness. Maybe tens or hundreds of probes or whatever.

We haven't really talked about what sort of locations in particular those things might be. Maybe they're in data centers. Maybe they're in exchange points. Maybe they're close to the servers. Maybe they're in, like RIPE ATLAS, in your house. We haven't really discussed that, so that's all up in the air.

We did talk a little bit about how, if you have this distributed platform, you probably need to account for anomalies. The probe needs to know if it's connected to the Internet or not so that you can throw [at] its measurements when things time out and things like that. So those sort of reasonable things.

That's about it. We don't have a lot more detail about how this might work. But your question about latency – is latency the right sort of metric? – my gut tells me that it is. It's the kind of thing that we talk

about in other contexts, like attacks. It's the kind of thing that DNSMON reports on. We can all go to DNSMON and we can look at the latency that it reports. We talk about it in RSSAC a lot when somebody says, "Hey, I want a root server in my country." We might say, "Well, you should understand how close or how far away from you before you really can answer that question."

So that's where I coming from on the latency issue. But, if people think it's not useful, then I'm happy to talk less about it or to take it off.

FRED BAKER:

Well, I'm not everybody. I'm me: Fred. And by the way, I think Wes and Jeff are wanting to get in a word in edgewise somewhere. I raised that question basically because – well, two things. One of them is I think that it would be valuable to have a downloadable instance of [inaudible] platform that can somebody sit down and say, "Do I need a new root server? Oh, jeez. The measurements are the same as [inaudible], so it's not a problem (or it is a problem)." I could imagine that factoring into the quality of service – whatever we called it [in] the work party – how bad it is, and where is it bad. They might to use that measurement platform to make that measurement.

But then also, if we're turning around and measuring a particular RSO, which is the other half of these – we apply similar metrics to individual RSOs – I'm not sure that one necessarily helps, simply because of the distribution of where some of us have our locations. So I'm not sure what it tells us.

DUANE WESSELS: So I understand that. I just want to point out that, in that concern, you're making the assumption, I think, that the probes are located at what I would consider end-user locations. So you're measuring wide area network latency. You're including that in the measurement.

FRED BAKER: That might be true.

DUANE WESSELS: Yeah, because, instead, if you did something the total opposite and you placed these measurement probes very close to the servers that are being operated, then you end up measuring server latency and things like that. You're not measuring network latency.

RUSS MUNDY: I think Jeff ...

DUANE WESSELS: Yeah.

WES HARDAKER: Thanks. Wes Hardaker, ISI. I guess we're going to talk about four metrics probably during the course of this week, and my guess that we'll come up with more.

With respect to latency, it is important these days because it's sort of the number one thing that everyone seems to be optimizing for, often

at the expense of other stuff, which I have a horrible issue with. So, eventually, we'll have to come up with some weighting factor as to which of these metrics is most critical to me. And I would say availability might be [inaudible] higher on the list than latency, for example.

But I suspect that we'll come up with multiple eventual metrics that should be met. One would be latency to be considered actually actively serving properly, and latency for a [tart] range. It's better than just a single value. It'd be more likely green, yellow, and red. If you look at the New gTLD Program, for example, I think the latency requirements in that contextual language SLAs is, like, 500 milliseconds, which could be hit from anywhere on the planet, including B from Beijing. We hit 200, I think, from Beijing, or something like that.

That's probably the right thing to think about: what's the latency to the service versus the latency to a single RSO? Duane and I received an e-mail message which I decided a problem shouldn't pass on the name because he wrote it just to the two of us. But I'm going to summarize the points that somebody else had sort of similar concern: that, by optimizing for latency alone, we would discard some of the other benefits that we get out of a diverse systems, such as different Anycast routing strategies and things like that that provide other benefits to the service. We do need to be careful that, if we optimize for a single metric, we will throw away some other valuable things.

UNIDENTIFIED MALE:

Are we using latency as a proxy, or are you thinking about – if this is the full set of measures for how tall you must be to ride, instant count and

geographic range have been two things that get used a lot when you're talking about "Are you tall enough for the ride?" But I don't see them here. I think that's part of why I was asking, "Is this a superset?" because, if they're not in there, they're not going to make that.

But if latency you're using as a proxy, then I guess I kind of get it. Sort of.

DUANE WESSELS:

Well, I don't consider latency being a proxy, but the work party has not discussed instance count or geography at this point. I probably should have said at the start that the work party is considering an RSO as a single thing and not trying to hit each instance or count how many instances there are. It's treating an RSO as a single service point. So I'm still not convinced this is the thing we're going to use for the bar or the SLAs then. I think there's a little bit of a mismatch. Just so we're aware of it.

RUSS MUNDY:

If I could inject a little bit, that is that, even though, yes, this is one of the metrics we're talking about, we have stayed at this point far, far away from any value – what is the required goodness value? – of any of the metrics. Is there a metric that can be done and used for this? So the value could be very broad.

DUANE WESSELS:

Paul?

PAUL VIXIE:

Paul Vixie, C-root. Two comments come up here. As the originator of the BPQ debacle, you have my apologies. But you're doing better with it than I would have. Latency is a proxy for all kinds of things, and certainly placement or locality is one of them. That's my first observation.

But it's also something that tells you quite a bit about buffer bloat or just buffering in general. If you have a queue of packets and bytes and things that are waiting to get into various network elements that are between the server and the requester, then seeing that float upward is an indication of load. It ends up being a proxy for how busy some pipe or some set of pipes is on your path.

I think, when you talk about latency, you are implicitly talking about that. So you might want to make it explicit and say that latency should not just be below some threshold but should have some small standard deviation. In other words, whatever it is would represent [queues] that are empty on average or very close to empty on average in all the network elements, from the server to the measurement point.

The second observation I would make is, when we started doing Anycast – M-root did it first, and then eventually F-root got around to doing it – somebody came up with the term “catchment,” as though rain fell everywhere but there was a slope to the ground that would cause rain drops to head west or east based on the topography of the routing tables.

So it's a very useful concept, and I think one of the things that I should have done back in those days and we should all do now is find a way to

measure from some idealized average query position within each catchment to see whether that service is performing well in that place.

Now, using something like RIPE ATLAS, you have pretty good statistical likelihood of representing every catchment, but that's not the same as being sure. If I were going to put a contractual parameter around it, I would say don't accept a so-called local node until you have at least two independent points that are inside of its normal catchment. Thank you.

UNIDENTIFIED MALE:

So two points. First of all, I agree with Paul that it should be deviation for latency, but the main thing is latency of RSS. We have to separate when we're measuring RSS versus RSOs. We shouldn't try, when we're measuring RSS, to answer questions about RSOs. So I think latency is a very good measure for the whole RSS, but you have to see RSS [bucks]. So what is that latency of getting to an answer for a root query?

For RSOs, that would be a very different story. I don't think we can implement that successfully, the reason being the routing world is much, much more complex and is run by many more parameters than just technical needs. I have something in Frankfurt which, with one hop, goes to Helsinki, but it takes 200 milliseconds. When we investigate, it's NTT's network, and this our network we routed via Japan and back to Helsinki because it's our own cables and we want to use them.

There is a lot of financial, like all this peering stuff and all of that, and it's over MPLS, so it's not invisible. And there's not much we can do. These tier-ones have their own rules, one of them here. And it's hard to argue

with them. Actually, most of them are not happy to discuss this. They're like, "Hey, my business. That's how I run it."

So discussing these things for with four independent RSOs and measuring things like latency – and also there are other limitations, like geo-satellites. It's 500 milliseconds, like it or not. That's the speed of light. Dialup is 90 milliseconds anyways, whatever you do. So on the technology of how the network is connected – the routing – there is a lot out of our control if you want to go [inaudible]

But for RSS, I think it's a meaningful measure. So we have to have a clear distinction between those two measurements.

DUANE WESSELS:

Okay. I guess I want to play ... well. I didn't know we were going to dive into the latency stuff so soon, but I guess I want to play devil's advocate on a couple of your points.

I feel like sometimes we're making assumptions about what latency thresholds might be in place. I [inaudible]. Was it Wes who said that the gTLD program has a 500-millisecond latency? So let's make it even worse. Let's say one second. Let's say our bar was that you had to respond within one second. Are you seriously telling me that you're still concerned about routing through Beijing and things like that that ... so you see what I'm saying? I think you may be having felt like, "In my [inaudible], I have to meet a 30-millisecond SLA," but we haven't discussed the threshold yet. We should be talking about what to measure and maybe how to measure it. And then we can talk about

high to set that bar. Maybe it's a really, really high bar or a low bar or whatever it is.

UNIDENTIFIED MALE:

Fair point. If I may comment on that, actually we have the same experience in ATLAS because the FCC apparently, in the U.S., has some rules for ISPs, minimums, which actually the bar is really low. But it's still [there]. Interestingly, many companies use ATLAS or many people want to complain to show to the FCC that they're not doing well. And we have been involved in some cases. Of course, we are based in the Netherlands, so we say, "This is neutral third-party data. Do whatever you want."

But as soon as you're in the middle of these kinds of contractual relationships, there's a lot of pressure. It doesn't matter what you said. People will try and find holds in every aspect of this to say, "Hey, this is not the right thing."

So I think, if you're going to base contracts on these SLAs, we have to make sure – I don't think actual numbers are the case. It's like the methodologies should be bulletproof.

DUANE WESSELS:

Right. And ... oh, go ahead, [inaudible].

STEVE SHENG:

There's someone in the queue. Paul Hoffman.

DUANE WESSELS: Go ahead, Paul.

PAUL HOFFMAN: Hi. Thank you. So I'd like to get off latency for a moment. I know that you folks will be discussing tomorrow. I probably won't be on that call. But I'm still confused about the relationship between the metrics that we're discussing here – the four that were listed plus the possible two additional of instance count and geography – and 037/038. Earlier, Brad said – and in the work statement it says that we should do these without 037/038. I don't think we discussed them outside of that context in the work party meeting so far, so why would anyone measure these? That is, why would any of these metrics be useful to anyone?

DUANE WESSELS: You mean why would they be useful if we didn't have RSSAC037/038, Paul?

PAUL HOFFMAN: Yes, exactly. So let's say that some researcher – let's not even say it's RSSAC – does some of these measurements and publishes them and maybe updates them. Of what value will those be, particularly because it doesn't seem like we have agreements yet on minimums or averages or anything like that.

So, really, what are the values of these, other than 037/038? That might help us move forwards a bit.

DUANE WESSELS: Steve, can you scroll this document down to the section where the RSOs send their responses? The summary? It's very near to the end.

Wes, you want to respond?

WES HARDAKER: Well, the only thing I was going to add is that the original point of this working group was to up ... well, that was fun. Yeah, it's coming out of ...

We're listening to him? ...

UNIDENTIFIED MALE: [inaudible]

UNIDENTIFIED MALE: [inaudible]

WES HARDAKER: Yeah. Can everyone on the phone mute?

Mute?

Paul, mute.

Paul is muted.

UNIDENTIFIED MALE: Test.

UNIDENTIFIED MALE: Yeah, same thing here.

UNIDENTIFIED MALE: Test. Yay!

WES HARDAKER: Now what was I saying?

UNIDENTIFIED MALE: There was feedback from their computer that their microphone was picking up from their speaker and sending back [inaudible].

WES HARDAKER: Okay. So hopefully we've got that sorted out, but I remember what I was saying. So the original point of this working group, even though we were talking about in terms of 037, was to update 001 because it's been a long time and update possibly the RFC into what should be used to measure the system and RSOs, regardless of whether 037 comes to light or not.

So, Paul, you're almost asking, is 001 helpful at all, right? I think people do look at it and [turn] to it, but I have no idea what people's opinions are of it.

BRAD VERD: Yeah. I think the output of the Metrics Group is kind of paramount. I think 001, if you go back and read it, is very vague. There's really no specificity in it, and there's certainly no teeth in it. We all agreed as a group that 001 needed to be updated, and we agreed on that over a year ago. Maybe two. I'm not sure.

So, Paul, I think defining what good looks like and providing that to the community is a good thing and a good start and would and should feed directly into 037 and 038 but are not dependent nor tied to 037 and 038, if that make sense. I mean, they're tied in that 037/038 would just use them. Does that make sense? It would use the output or the metrics. But I feel like we're getting lost in this, but – Jeff's looking at me funny.

JEFF OSBORN: Well, if we can say later that the things we're using to measure the bar or anything like that can be added onto the work they're doing, then it's no big deal. If the work they're doing is the sum total of anything you'll ever get a hand on, then it feels like whoa, whoa, whoa.

BRAD VERD: Stop being an engineer for two seconds. Nothing we have ever done with DNS is a sum total at that moment. It's always evolving.

JEFF OSBORN: Okay.

BRAD VERD: And we're constantly changing it. That should be the expectation here. Everything we do, every document – we've updated our procedures document five times now, right? So this is constantly evolving and will evolve, just like we added DNSSEC in 2010.

DUANE WESSELS: So I asked Steve to put up this section of the document. This is a summary of responses that we got back from root server operators. If you scroll down a little bit, Steve, I notice that some of the responded that they already measure latency. They already look for this. They already do both DNS response latency, and some even do ICMP latency. Now, I don't know how many of the operators that responded do that, but some of them do.

So it seems to me that that we're already doing this to some extent. This work party is about making a definition that we can all agree on as to how you would measure this across the board.

I think – well ...

BRAD VERD: If I may, just a couple comments I wanted to add before we get sidetracked with the phone there. One is I hope that we're not engineering to the 100% because I don't think we can find a solution that will cover every instance, meaning we shouldn't be engineering for the one-offs in China or the one-offs in Antarctica or whatever. There's going to be nuances because of topology is not black and white. It's different. Just like you were saying, NTT sends their traffic back to China

and back over their own cables. We don't have a solution for that. It's kind of like the 80/20 rule. Engineer to the 80 and hopefully that'll give us a good enough answer. Maybe there's a Venn diagram in there, Jeff, but we get a good enough answer. I'm not sure.

Going back to a couple comments that were said earlier, I feel that, if you talk latency or measurements for the service, those levels should be lower than they are to the RSOs because the service that's provided by all twelve organizations should be in theory more distributed, more available, more, more, more, more. So those numbers should be lowers. Then, if you take the same numbers, and apply them to the RSOs, there's going to be a bigger spread or a bigger average that maybe you guys mentioned earlier.

So I guess I say that because I don't want people it's the same [spaim.] I see them as different. When you talk latency, I think latency becomes a proxy for a lot of things. I think latency could be a proxy for instances in geographic location. If you look at latency to an individual RSO, if latency is super high, maybe it's because they've only got one or two sites, where, if you move it – I mean, we had a study that we shared with the RSOs a couple years ago in Prague when we took A-root from ... what was it? I think we took A-root from—

UNIDENTIFIED MALE: Four sites.

BRAD VERD: Four sites to eight sites? Something – it wasn’t dramatic. It was just a couple – and we literally exponentially made it ten times more [available], or the latency dropped dramatically for everybody around the planet. It wasn’t like you have to have 100 sites – there was this ROI that drops off when you reach a certain number.

But anyway, those are just things to think about. This is not going to be black and white as we talk through this. That’s all I’m trying to get at. And that’s going to be hard for some of us to wrap our heads around.

DUANE WESSELS: So I wanted to see if Paul Hoffman is back on the call and if we got his question answered and if he heard the discussion or not. Paul, are you with us?

PAUL HOFFMAN: I am, and I’ve heard the whole thing. I just typed a brief thing in the chat, saying that the work statement that the work party has been working from actually doesn’t mention either the word “good” or “sufficient.” So, since it sounds like those are things that are now being discussed much more by RSSAC, when it comes back to the work party we might have a better target.

DUANE WESSELS: Okay. Thanks, Paul. I think it’s important to keep those in mind. I don’t know if we necessarily need to go back and update the statement of work. But certainly I think that’s a good direction to stay in.

LARS-JOHAN LIMAN: I think this work should remain focused on what can be measured and what impact does the different variables have. Setting the actual bar I see as the second thing that we may or may not want to save until the bigger discussion on where this is heading; the entire revolution of the system. But to first find out what can be measured and what's the impact of the different variables is a very important first step.

WES HARDAKER: Yeah. As a shepherd, I think that we intended to define what metrics were, not good and bad. I think we very deliberately stayed away from that, and it's natural in this group when we are talking – and we just finished talking about 037. To want to dive into how do we measure what's good and bad? But I think if we could get through defining what to measure first, and recognizing that there's some value. But we shouldn't set the values now for bar levels. That's out of scope.

DUANE WESSELS: Okay. I'm okay with that, but—

UNIDENTIFIED MALE: [inaudible] setting—

WES HARDAKER: Setting bar levels today in this document. This document is – yeah – is not supposed to be producing values.

DUANE WESSELS: Let me respond. So, in the document that we have prepared that we're going to go through, you will see a lot of blanks where a value could be filled in. We don't have to fill those in right now. We could say that's future work for the work party or that's future work for another work party. Whatever. But just to prepare you that that is in that document.

BRAD VERD: I'm sorry. Of the four deliverables in the statement of work, one of them was to define what good looks like, right?

LARS-JOHAN LIMAN: No.

BRAD VERD: I'm sorry. I missed – maybe I'm ... You had it up. I don't have it up in front of me, but—

WES HARDAKER: No. As actually somebody else pointed out, the word is to define what the measurements should be, not to define good and what you must pass of a measurement. It's to create the measurements, not to define a particular value of a measurement was good or bad.

DUANE WESSELS: So I would suggest that we save that part for the end. Let's focus on the metrics themselves. If, at some point, RSSAC or the work party decides it has the appetite for setting those thresholds, we can do that. But it shouldn't be the priority.

WES HARDAKER: Yeah. Certainly we shouldn't stop conversation, but that is future work that has to be done. So there's no harm in continuing to discuss it, but as far being this particular work item done, it has to be about the ...

BRAD VERD: So I pulled it up in front of me. I'm reading it. Maybe I inferred it. Let me ask the question differently. What do you see as minimum level of performance?

DUANE WESSELS: Well, it's the second bullet on the screen. Niall and I were just talking and was going to make the argument that I think "define measurements" implies to metrics and then a measurement for those metrics. The way I read that means the big picture is to find out what's good enough, what's tall enough.

BRAD VERD: We can change this if we don't want to do this, but the goal here was to define the technical accountability. So the root server system and the root server operators. So, if we are going to say where the bar is, that I feel is a failure on our part.

Paul?

PAUL VIXIE:

I agree with you. This all started with RFC 2010, which was written by a couple of root ops. A lot of people looked at it and said, "Yeah, but you forgot about all those other stuff," and it seems relatively self-serving for you guys to define what goodness is for root ops. The idea of having the [Caucus] do this and have it be done with a lot of public oversight and transparency is really so that we can move beyond that and get to an objective measurable standard about whether you're doing this well enough to properly serve the community or whether you have somehow fallen out along the way and can't catch up.

So I think the purpose of this is the same as what the purpose was of RFC 2010, and it isn't just to feed into RSSAC 037 and 038. It is so that the community can determine for itself how tall you have to be to ride this ride. We're counting on their trust. This is the way they will know whether we are earning it.

LARS-JOHAN LIMAN:

I understand that wish, and I actually support it. But I think we're stepping ahead of ourselves because this is what I envisage – which function would it be [that] the SAPF function in the future model would deal with? The mental picture I had of the current work here was to figure out which knobs are there to turn and which dials are there to look at and then for the SAPF to have that display [inaudible] "Okay. We want the dials to be at least this high, so the green sector on the dials is this." If you want to define that already at this stage, then we could do

that. But then it's a different thing for me because then we're setting the bar for the current system. Then that may be a good thing as well. But it's a different thing for me. And that's fine. If we want to do that, that's fine.

BRAD VERD:

So look. I've said this a dozen times. I'll say it a dozen more. I will get up on the table and dance it if I have to. And that is we should not be defining anything with the current system. We should be defining what good looks like. I would argue that maybe [inaudible] today isn't what good looks like.

So I hope that's clear. I don't want anybody to walk away from here and be like, "Oh, Brad said today is what good looks like." I'm not saying that. I'm saying, if we were to start a new slate, what would good like? What should we measure?

LARS-JOHAN LIMAN:

Fair enough. It was expressed from my side. The difference I want to express is between looking at the point in time now and the point in time where we have a future governance model in place, operating. And that—

[BRAD VERD]:

What changes?

LARS-JOHAN LIMAN: What changes is who gets to decide what because in the future model we have the SAPF and we have other functions that we don't have today. They will play a role in this.

BRAD VERD: They could update the document. They could take stuff away from the document. They could add stuff to the document. But why can't we define what we think good looks like? It's the community. It's the same expertise. What's going to change? If 037 were to go in tomorrow, a year from now, or four years from now, what would change from today for us to define this? I have a hard time understanding what's going to change, what magically is going to happen, when 037 gets implemented.

LARS-JOHAN LIMAN: You get possibly other and more people at the table who—

BRAD VERD: They could update it.

LARS-JOHAN LIMAN: Yes, but then back to my question. This was my mental image. What are we deciding for here? The time gap from now until the new system is in place, or are we trying to [design] for the new system or laying the ground for them or something. I hear the former, and, okay, I buy that. So that's fine.

BRAD VERD: Again, I'm sorry if I – I feel like I've been so clear on this. In Barcelona, we said, what work could we pick up in lieu of 037 taking a long time to happen? What could we take on and get done that would just plug in to whatever model if it got implemented, or even if it didn't get implemented. What happens if 037 doesn't happen? Three years from now, all of a sudden we still haven't defined what good looks like? That feels like a real failure.

LARS-JOHAN LIMAN: Fair enough.

BRAD VERD: So, to me, technical accountability is something that is not dependent upon the governance model. The governance model should be able to use it. It should be able to plug right into it. Then, whatever the model is, they can update it. They can change it. They can throw it away. But it feels like we have the right group, the right expertise – we have the community, we have the transparency in place – to do this. So I'm asking the question, why do we not want to? Because I feel like there's pushback here to not want to do it, and I don't want to understand that.

KAVEH RANJBAR: May I—

BRAD VERD: [inaudible] from Kaveh.

KAVEH RANJBAR: Oh, fine. No, In defense of trying to actually come up with what we feel is good from scratch and not looking into what we have to today, there's also a risk – which also came up in the morning presentation by Jeff but then it was edited in the document – which is there is a case for going back to the average. We say, "Okay. This is what we do, so that's average." Average always means there is something below it. And that would cause issues. If you want consensus and if you want to go forward, I think we should actually say, "Okay. We are designing a system. What are the metrics? How do we know this system is a good system?" Talking about today, we all might be below or we might be higher than it. But I think it's wise to not try – we should learn from what we've done, but we shouldn't try to base it on, at least if we can, the current situation.

DUANE WESSELS: Andrew? Oh, Daniel Migault has joined online. Welcome, Daniel.

WES HARDAKER: I think a lot of people had their hand up. I know mine was one of them. So let me rephrase because I don't think anyone is saying we shouldn't do that, Brad. I think everybody wants to define the bar. That's not the problem. It's a question of then ordering.

So we have two choices, two paths forward. Honestly, I don't care which direction the group goes. But we can define just the metrics, not

how high you have to be but whether we should use inches or standing stones in order to figure out whether the alien from Mars is tall enough to ride the ride and then later we define how many stones or how many inches you actually have to be, or we define them both at once.

There's value in both. The problem with doing them at once is you end up with a boiling-the-ocean kind of case, but at the same time, if you don't consider numbers, then you might define metrics that actually aren't usable.

So we can errors in either direction. I'm fine with either direction we want to go. We have a lot of time here to talk about. I personally tend to believe we ought to – I would lean toward wanting to define the metrics without actually putting in numbers for good, and then doing good immediately thereafter.

[KEN RENARD]:

I definitely agree with that. I think some of the metrics, specifically latency, and if you're putting a number of milliseconds on it, it depends on how we measure. So if we did what are the metrics, maybe start to define how we measure, and then put final numbers on it. We can put numbers in just for the sake of discussion, but we'll probably go revisit the actual numbers at the very end.

RUSS MUNDY:

So if I could interject here as one of the Co-Chairs, I think we may be getting a little over-exercised, possibly, on whether or not the work party is supposed to be defining what's good enough. When I read the

second bullet, I as one of the Co-Chairs say that that's part of the job of the work party.

But, so far, what the work party has been elected to do more or less follows what Wes just suggested: first identify what are the things themselves that can be measured as it's laid out – in other words, get a better handle on what the metrics should be – and then look at what are the measurements associated with metrics and how to take them [inaudible] so forth. Then, thirdly, what are the thresholds? How bad is bad before something happens?

So I think the work party so far as been following the suggestions that Wes just laid out. So do you agree?

[DUANE WESSELS]: [Do you want one of the last words?]

LARS-JOHAN LIMAN: No. Just for the record, Brad has convinced me and I probably have to [inaudible] has convinced me again.

DUANE WESSELS: All right. I think it's time for a break. Am I correct that we're a few minutes over the allotted time? So how long is the break?

15 minute break. So we'll back here at about [2:50] or so. Be prepared to talk about BPQ. All right? Thanks.

RUSS MUNDY: I think break time is over.

UNIDENTIFIED MALE: [inaudible]

[DUANE WESSELS]: Okay, thanks, [Brad]. I could see where you'd ask that.

STEVE SHENG: Okay. Over to you, Duane and Russ.

RUSS MUNDY: Okay. So, could we put the document back up in the packet bandwidth queries section again? Since this session, it still says the leaders are Brad and Fred, but I think they were pointing our way.

UNIDENTIFIED MALE: You were looking very hard.

RUSS MUNDY: I thought so, yeah. I think we probably just didn't get the people names changed with some of the sessions shifted around here. We have folks on the call again. Paul was with us earlier. Do we have other people on the call? Do we know?

[ANDREW MCCONACHIE]: We have Kazunori Fujiwara and Daniel Migault.

RUSS MUNDY: Oh, great. Great. Thank you, Andrew, for keeping eye on the room. Just raise your hand to interrupt or whatever if we don't catch you.

So this section is just simply labeled BPQ. Well, BPQ is in the work party statement of work. In fact, we've had several discussions in the work party around BPQ. One of the few conclusions that we've reached is we don't know how to reach a conclusion, particularly for this work item.

The part of 037 that makes reference to it Duane has nicely put into the document that's on the screen here. This is in the Zoom room also. Is that correct?

UNIDENTIFIED MALE: They have the link to that.

RUSS MUNDY: Great. Thank you. So what we've got is the concept that is described in SAC037 of BPQ that came out of the very fruitful of our friend, Paul Vixie, as he admitted to earlier, and was a big help in completing RSSAC 037 and answering questions that we, as RSSAC, were being pushed on to answer, particularly with respect to cost from the Board and from the CEO at the time.

So the section that addresses this, in 037, I extracted here, just so folks have a fairly consistent context. In the discussions that we've had

in the work party, as I best recall them, when we tried to discuss it for the RSS system, there were many that raised many concerns with respect to even identifying that information because, in a way, it would give a fairly visible, quantitative target for attackers to go after if such a thing did get defined for the system. They just know, “Oh, I just need to raise my attack machinery to blah.”

So that was one of the early things that was identified, but as we got further into the discussions, I don’t know that anyone really was able to come up with an approach for applying BPQ to the system.

Now, for individual RSOs, I think people in general felt that it probably could be defined and we could come up with metrics for it, but the same sort of concern was raised with respect to being a very obvious target to exceed. I can take Q-root [off there] by exceeding its BPQ. So that was put out a couple of times by folks.

One of the concerns that also was raised is that, over time, yes, maybe it’s a good financial measure, but people – I think in the feedback that we got – were saying, “I don’t know that it’s that good of a performance kind of metric.”

So, generally, we’ve had a lot of struggles trying to make any useful advances in terms of the Metrics Work Party defining what it would mean or how you’d go about defining it for either an individual RSO in a way that would comply with the other things – be public, be measurable, etc., by everybody.

Brad?

BRAD VERD:

I really hate being put in a position to defend the BPQ stuff, but I'm going to ask a few questions. One is BPQ was never intended to be a performance metric, I think. It was put in place to answer a very specific question of how much it costs. The reason – again, if I'm out of line here, please speak up, Russ or anybody – it's here is in this Metrics Work Party is, if you define what good looks like (if we get there), I think we've all stated that the very next question is, how much does it cost?

So, if you define what good looks like and there is some costing model, then that would answer that question. That's why it's here. If it can't be here, then it can't be here. Then that's fine. It was never meant to be a performance metric. If we wanted to make it one, cool. I don't know how you'd do that, but that'd be neat. But that's the genesis of the BPQ stuff, as I recall.

[PAUL VIXIE]:

I think that is why it is in 037. I do not think that is why it was brought into the discussion that led to 037. It was not meant to be a goodness metric. In that sense, it is not a performance metric. It is, however, a capacity metric. The specific backdrop of mentioning it was, depending on what your kind of routers and switches and load balancers and firewalls and line cards and hosts and operating systems and [virtualizer/hypervisor [inaudible], there's a whole string of stuff upstream of you between you and the Internet core, and it may be that it can handle an awful lot of bytes, as long as the packets are really large, or an awful lot of packets as long as they aren't fragments.

But, ultimately, it comes down to those three numbers. You have to exceed one of those three capacities to have a capacity-related service outage on that node because, no matter what you're doing, no matter how big the packets are or whether they're fragmented or they've got TCP, there's some number of bytes above which you can't cope, some number of packets above which you can't cope, and some number of queries above which you cannot cope.

So that what it was brought in for. I realize that it has been repurposed because it was the only thing nearby to the problem we had. But just to be clear, it was never intended as a metric of goodness. It is a metric of badness, the threshold above which the red light comes on. It is not an indication of where the green light should be. Thank you.

RUSS MUNDY: Others with thoughts on BPQ? Ken?

KEN RENARD: It inherently is a measure of capacity. Maybe could we turn this on its side and have it a measure of performance? And do we want to do such? Maybe you could take the current load, multiple it by 120% or something like that, and that's goodness. Is that something worth even doing? I'm just trying to think of a way that we could turn this into how good is it or a measurement. I think it's a good idea.

RUSS MUNDY: Fred?

FRED BAKER: Well, I struggle with this in that context, and I struggle with the idea that this would tell somebody how to attack my system. If I've got 50 or 100 or 200 nodes, you're going to have to attack each node individually in order to take the services offline. Setting up B times [P] times Q in one place in one place is not going to hit all the nodes. So that just seems a little simplistic to say that that's giving somebody now testament for how to take the system down. I don't think it is.

I also worry about the difference between peak and average. If we're reporting daily how many packets are going to all of the nodes that some operator is doing, which is what we do, that doesn't tell me anything about a particular five-minute interval or a particular second, and it doesn't tell me anything about a particular location.

I just struggle with that as anything other than measure of bulk capacity. I don't know how to interpret it.

RUSS MUNDY: Duane?

DUANE WESSELS: So, in the previous section, when were talking about latency and availability of things – are you just in the queue, or do you need a [inaudible]?

UNIDENTIFIED MALE: Paul.

RUSS MUNDY: Paul, go ahead.

PAUL HOFFMAN: Okay. Actually, I don't want to disrupt the technical flow here, but one thing that wasn't mentioned was that, in both in the last call and in the face-to-face meeting, people in the work party pointed out that the work party is tasked with coming up with externally-verifiable measurements. And none of BPQ are actually externally-verifiable. It doesn't mean we should have them. It doesn't mean you want to change it. But just a reason that we were having difficulty with BPQ was that they are measured by the operator, not externally-verifiable. Thanks.

DUANE WESSELS: Yeah, that's right. That is one of the, I think, struggles: this is different than the other things that we've been taking about. The latency and the other things are all designed to be externally-verifiable, and this one is just different. So that was one of the points.

Go ahead, Brad.

BRAD VERD: I guess I don't connect the two. BPQ doesn't have to be externally-verifiable if it's a costing mechanism. If it's a costing mechanism, it's a

costing mechanism. But I don't know if it works as one. I'm just saying that was what it was presented as in 037. So, if you define the metrics that are externally verifiable and you've got – I don't know – some sort of latency and there's ... yeah. I don't tie them together as externally-verifiable. So maybe the question is not BPQ but how do you cost this?

DUANE WESSELS:

Well, that may be the question, but that seems very different than an RSSAC Metrics Work Party task to me.

RUSS MUNDY:

Yeah. And let me just say, as one of the Co-Chairs, Duane and I have talked about this about, and that is that I think we were both trying to get our heads around how it is a metric. And of course, thinking of it, how can it be a performance metric if it's truly to be looked at as a costing metric? Then I would suggest that it is fundamentally different than the other pieces of the metrics that the work party has been asked to identify.

We can certainly still try to do it, but I think most of the people that have been involved in the discussions have been trying to get their heads around how is this performance. And that's been part of the challenge.

KEN RENARD:

Another way to think of it even beyond costing is we talked about the possibility of bringing on a new RSO. We can measure latency. We can measure correctness. Things like that. But wouldn't we want to measure

a new RSO's capacity at some level? It doesn't have to be their max capacity, just – at least we could probably define a minimum bar threshold that a new operator might have to meet, but their actual capacity could be order of magnitude beyond that.

RUSS MUNDY: So you're suggesting, Ken, that this is an appropriate thing for the Metrics Work Party to handle?

KEN RENARD: I'm just suggesting a new way to think about it [inaudible].

DUANE WESSELS: So, following up a little bit, Ken, do you think that would be something that would be actually measured, like with live traffic, or it is just, "Tell us what your capacity is and we believe you"?

KEN RENARD: In the context of a new operator that's not online in the [inaudible] yet, you could. You could actually measure it. But, yeah, it becomes impractical sometimes to do that.

[PAUL VIXIE]: So most of us have a standard model that we deploy such that – we may have more than one. You've got the newest one. You've got the slightly older ones. You've got the really old ones. But nevertheless, for any given Anycast node in any of our networks, we kind of have an idea of

what it is – what the model is, how much memory it has; that kind of stuff – and asking each one of us to please put one instance of that in a test lab and beat the hell out of it and show our test results? We wouldn't have to get an external auditor to do it for us. We wouldn't have to ask OARC to do it for us. We're just saying, "Please show us that you have beaten the hell out of this and what you got as a result, and then we'll check the box to say that you have reported your BPQ." That's a little from, "Just wave your hands and estimate what you think it probably is," and it's a lot different from proving it to an outside party. I think there's a very reasonable middle ground now.

DUANE WESSELS:

Paul, I'm going to make sure – so you said something like showing ourselves that we can do it? But you're not proposing a third-party be involved in auditing or measuring or ...

PAUL VIXIE:

I think this is almost like trust but verify. I'm not asking that, let's say, RIPE submit to some kind of an annual audit of one each of the different kinds of configurations that they are using for K-root. I'm saying, if they could, every time they are going to deploy a new style of thing – a new major version of the OS or getting rid of the firewall or building that into the host; whatever it is; we're upgrading, we're using a better class of rack now, PC now, than we used to do – self-certify, just say, "We tested it in and this is where it broke: In B's, in P's, and in Q's." That, I think, would be so much better than we have now that we wouldn't

necessarily ask them to prove it with a CPA in the room or a lawyer or an accountant.

DUANE WESSELS: So one thing we don't have now is any kind of a tool that anyone knows how to use that can – I mean, there's maybe a couple. But do you think the work party should define how that test looks and runs and how long to run it for and what kind of traffic and that kind of stuff? Is that ...

PAUL VIXIE: If the work party has an appetite for that, it would be of great value. Generally speaking, the tools that Nominum has given away, and Netperf 3 are enough to do this. It's just matter of having a standard recipe of how you configure them and how you run them and how you interpret or maybe even graph the results.

KAVEH RANJBAR: Clarifying question because, when we define BPQ and when we present it as it is in 037, it's for the whole RSS. So we never discuss that per RSO – I mean, maybe [inaudible] to collect it for the RSS. We do individually and then sum it up somehow. That I don't know. But the idea is not to measure per RSO but only for the whole system. And we should come up with the methodology. But the goal was for the whole RSS, and that's also what the tech says.

DUANE WESSELS: Which makes it a lot harder.

KAVEH RANJBAR: Right.

DUANE WESSELS: Somebody—

UNIDENTIFIED MALE: [inaudible]

UNIDENTIFIED MALE: Daniel Migault comments, “Should we define BPQ on a root server instance, i.e. a node? Capacity of the RSO could then be extrapolated from the performance of a node.” Then he adds, “You can add that it takes a [monola’s] input. I agree that this may be outside the scope of the Metrics Work Party.”

KAVEH RANJBAR: So as a reminder, just to answer that comment, BPQ [is] as it is in 037 for that purpose. Maybe we want to use it for other purposes, but the idea was to give – we basically asked the Board ... The discussion was, “What is your comfort level with the performance of the root server system?” We came up with BPQ and we decided that that’s our proposal, that we would give you a number for the current BPQ that the system can handle, whatever it is, and we never got into the methodology. Then you can say, “We’re happy with that” or, “We want to have ten times more than that,” or, “a hundred times more than

that.” That would give us an indication for – not us, the function; SAPF function and others – on how to plan adding new ones or spending money on improving what you have today or anything.

So it’s basically a comparison measurement. It’s not designed for the health of the system. It’s just basically to say, “Okay, we’re at this level now, and the ICANN Board thinks it’s fine,” or whoever. I’m just using the ICANN Board as a decision-making body. But, “We are fine with that [risk level],” or, “We want ten times more than that capacity and then we are fine.” That was the main reason that we came up with this, in the context of 037. Maybe we want to use it in other places for different reasons. That’s fine. But I just wanted to, as a reminder, tell you how it came to be and what was the purpose of putting in 037.

RUSS MUNDY:

Okay. Thanks. Kaveh. So we’ve had several suggestions here for how we might handle BPQ in this context. My original personal thought coming into this was it doesn’t really fit with the rest of the tasking for the Metrics Work Party and maybe should be withdrawn from the SOW. But we certainly have had a number of other suggestions, and that one did not come up; of withdrawing it from the Metrics Work Party.

So I don’t know if we can have a little discussion about what I’ll describe as Paul’s suggestion/approach, which is to come up with some type of mechanism where each RSO would do a self-test, self-certification, against some level using existing tools. Then also Kaveh’s suggestion was in a way – why 037 said what it did, which essentially was saying we need to figure out a way to say BPQ or the equivalent for the entire RSS.

So I'd love to hear inputs from people on that because I don't know how to do what you suggested, Kaveh. I think Paul's suggestion is quite an achievable thing that could be fleshed out and written down. But we're in extended work party discussion here, so this really is viewed as part of the work party.

Open for comments from the floor. Thoughts from people on how and what we ought to be doing with BPQ?

KAVEH RANJBAR:

Question for people who are familiar with 002. I think, for each individual operator, we have actually BPQ as part of 002 because we measure packets, we measure queries, and I think we even bandwidth [inaudible]. I'm not sure. So, for the current state of current load, I don't think it's that hard. But if we want to go current capacity, either of the system or each operator, then I don't know.

DUANE WESSELS:

So you could envision some kind of reporting [on] RSSAC002 that gets to BPQ, I think. You have to make some assumptions. But you could do something. It would be at the granularity of one day, which I forget who – Fred had made the point earlier about averages and peaks. So that's a pretty big interval. But you could do that. Yeah, I think you could do that from RSSAC002 with those [caveats].

RUSS MUNDY:

So could we have thoughts from people that are familiar with the collection of 002 right now?

DUANE WESSELS: Can I ask a load question, I guess? Let's say that we did do that. Let's say that somebody went to the trouble of analyzing the 002 data and said, "Okay. The root server system as a whole – the current load is X bandwidth." Are people in this room comfortable with that data being published and out there? That doesn't ... because—

PAUL VIXIE: RSSAC002 data?

DUANE WESSELS: No. Let's say that there was a report that said, on March 31st, the root server system handled X bytes per seconds, so many queries, and so many packets.

PAUL VIXIE: It takes a small matter of programming, but that stuff is online.

DUANE WESSELS: I know. I'm saying we can do it. I'm saying, would you be okay with those numbers being published? Because I've heard from many people that they do not their capacity numbers published.

PAUL VIXIE: Wait, you just said—

DUANE WESSELS: I'm talking about two different things, but are you okay with current load being published while not being okay with capacity being published?

PAUL VIXIE: When I used to have to make this decision, I never published instantaneous load because 400-lb. teenagers living in Trump's basement are very interested in bragging about what's happening right now. And if they have to wait 15 minutes before they can claim victory, they will move, and they will find some target who has more real-time results than I do. I still feel that way. But with a 15-minute delay, I have no problem with publishing current load.

DUANE WESSELS: Well, with what we're currently discussing, it would be from RSSAC002 data, which covers a whole day. So that would be a one-day value and it would be delayed by, probably, one or two days.

PAUL VIXIE: I have no objection.

DUANE WESSELS: [inaudible]. Okay.

PAUL VIXIE: But I want to answer the second part of your question having to do with capacity. If we tried to reason from the measured maximum B,

maximum P, and maximum Q of every Anycast node and we could somehow keep that information up to date as they move around and as they come and go and try to say that the aggregate capacity of the system was computable from that, it would be laughable. So, other than the danger of being laughed at and losing our credibility, I have no problem with publishing those capacity numbers, imaginary though they might be.

RUSS MUNDY:

So, Kaveh, a question about, as you stated, the original purpose. It almost sounds like – if we took an approach that said we'll compute our daily load numbers and publish that as a BPQ, would that handle the problem that you were describing earlier about how big is this system and how much risk are you as a Board willing to take as so forth?

KAVEH RANJBAR:

I think the idea was more for the capacity than the current load. The current load, I think, can be useful for us and to set the baseline or to show, okay, this is the current load. But the important thing is capacity. And actually showing those two different is very important as well because we can say, "Hey, this is the current load. This is the capacity we think we can handle," and then, "Are you happy with that? Or do you want ten times more than that?" or less or whatever.

But I think for decision-making bodies, capacity is important. I assume.

RUSS MUNDY: So I interpret your answer to say that I would be useful information to do the BPQ for load on a daily basis. But it wouldn't really answer the question the Board was asking. Is that correct?

KAVEH RANJBAR: Yeah.

RUSS MUNDY: That's what I thought. Okay. Other suggestions on how we might give that answer?

KAVEH RANJBAR: I think what Paul suggested is a doable, workable solution. It's just like we come up with the numbers. How I guess is up to each operator. For individual operators – because I know, actually, to some degree how much capacity I have over the current load. So I have the current load I can extrapolate and say, "Hey, five times more than that? I'm sure I'm safe. We have even done tests, blah, blah." So for that ICANN provider – I don't know if other operators are happy to do so – I think that would be a possibility.

And the other question would be I don't know if we can simply sum them up or that would be a different – I think that's something we need to investigate further.

BRAD VERD: I'm sorry to keep bringing this up. Maybe I misunderstood BPQ entirely. For BPQ in 037 – I know that the Metrics Work Party is not tied to 037; this is where it started, so that's why I go back to – we were asked by the Board to put a dollar figure in. We didn't want to put a dollar figure in, so we tried to put a model in instead. So we put in BPQ, which basically said, "If you can figure out the bandwidth, the packets, and the queries, you should be able to go off and take industry-standard numbers and model what the system should look like to handle that BPQ." Am I correct so far?

UNIDENTIFIED MALE: Yeah.

BRAD VERD: So then, if you come up with that dollar figure – let's say it's \$100 – to run the system today, then came the risk factor that we said had to be determined by the Board. So the Board said, "We want 100X." It was 100 times whatever the current dollar figure is. So 100 times 100 – what is that? 10,000. And that would be kind of the mode.

How is that not doable if you have the numbers from RSSAC002?

KAVEH RANJBAR: I think the difference is between the capacity and the load. Correct? So the load might not be that much – I don't know. I'm just trying to string it together. So if I know the current load but I don't know how much capacity you can handle, that's not much useful for that extrapolation because your capacity might already be 100 times more than what

you're handling. Then, if I say, "Oh, I'm happy with ten mores than that," that would be a kind of ...

BRAD VERD: But going to Jeff's comment, there's probably twelve different models for capacity. Is that a fair assessment?

[KAVEH RANJBAR]: I think so, yeah.

BRAD VERD: All right. So, if there's twelve different models for capacity, it's really a question for the Board that is, "If this is your normal capacity" – and we have all the stats on that; we have years of stats on that – "and here's your [growth rate], what do you want your overhead to be?" Then they say ten. They say four. They say five. They come back to us. We say, "You're good." Or they come back to us and we say, "Well, five of us are good. Three of us are on the border and the other four are below. So we need money."

KAVEH RANJBAR: Fair point. I think that's doable except it's for the whole system, not for individual operators Because the idea was not to involve them in that discussion. It's basically like the whole system has this load, if you want to say, and you say you want to be responsible for – you're feeling safe if you think we can handle 100 times more than that; the current load. Then we should decide – or the functions – how we can provide that

kind of capacity. Then, how you translate those individual roots or say, “Hey, one can be less, but there’s” – that would be all up to us, not the ICANN Board. So the idea was to provide BPQ only for the RSS and do the comparisons or [ideally] only for the RSS, not the individual [roots].

BRAD VERD: I understand that. I’m going one step further to say it now becomes a tool used by this group to justify funding, should it be needed, especially if we define what good looks like, which ties back to here; ties back to the Metrics Work Party. If the Board is saying that 10X is what good looks like—

KAVEH RANJBAR: I think that can work. We can also work on load. [You] are right.

RUSS MUNDY: Daniel, go ahead.

DANIEL MIGAULT: Just back to – can you hear me?

RUSS MUNDY: Yes.

DANIEL MIGAULT: Okay. So back to what Kaveh said, he said, “Well, I’m confident, given a load, that the capacity of my system is five times the load. It’s fine.” I

think there is a huge assumption, which is that the capacity is measured in terms of load. Just because that supports five times the load doesn't mean your system is resilient to an attack that is equivalent to five times the load, especially if you target that locally on some area or at some specific time and so on.

So I think it seems that, if everyone is rezoning based on the load, it might help. Then you may assume that, if you're able to address ten times the load, then you might be able to address the potential attacks. That's a decision the Board may say.

But taking the unit as the current load might help us also.

PAUL VIXIE:

The assumption that I make – if I'm right to boil it down to bytes, packets, and queries, and if I'm correct that we are capable of reliably and honestly measuring each of various configurations that we have in network to say this is where the red light will come on at each of those is exceeded, then what I would say is, whatever the current load is, represented as a fraction of that threshold where the red light comes on, one of them is likely to be higher than the others. You might be at 50% of your capacity for bytes just with your daily capacity. But you might be only at 10% of your capacity for packets and queries, or vice versa. But whichever fraction is the largest is, to me, the one that is most in danger.

We're throwing rocks to try to hit the side of a barn here. This is not an exact science. But if none of your measured loads in those three metrics comes within 20% of your measured capacity, then you would be not

irrational or dishonest to say that you have 5X headroom or 4X headroom; that you can stand five times as much traffic as you're seeing. It doesn't really matter, since you chose the largest fraction, which of the numbers rises first. It will still have to be at least five times as big as your current load in order to make the red light come on.

So, yes, I recognize that this is bad science and that I would be tossed out of math class for thinking this way, but it's actually way better than what most of the industry does.

RUSS MUNDY:

Anybody else online – oh, okay. We've heard from several of the current people that are quite familiar with the 002 stats and handling, and we can derive a day-based period for a BPQ number from our current 002 stats. I'd like to just ask that open question to any of the people familiar with their particular root operation if they see any problem at all with this, or is this something that the work party can take and proceed forward with – this building on this approach – at least to begin with[?] 002 stats, which have a day-period length to indicate load, can be used to derive BPQ. Does anyone see a problem with that?

Great. Okay. That is itself, I think, a big step forward for us for what the work party can do with BPQ. I think we've got a better, useful direction here for that.

The other piece that I think we've generally agreed to – but I'm not sure of it yet, so let me ask this. I didn't hear any discussion that indicated that we, as the work party, should try to do any kind of BPQ for the system, for the service itself. And—

UNIDENTIFIED MALE: Do you mean separately form what you just—

RUSS MUNDY: Separately from doing it for the individual RSOs. This could have been part of me. I don't know if Duane read it this way, too. I read the charter to say we were supposed to do it for both the RSOs and for the system. But maybe I was misinterpreting that. So do we need to do it for the system?

BRAD VERD: Sorry. The BPQ thing that you just recapped I thought was for the system. Now you're saying it was on a per-RSO basis?

RUSS MUNDY: Well, [we] I think have got full agreement from anybody that's spoken up for it [that] each RSO is okay with the BPQ based on the 002 stats for each RSO. We hadn't really talked about how one would go about doing a combining of that in some manner to reach a BPQ for the system. So that's really the question, Brad. I think we have agreement for the RSO level. I don't think we've reached agreement as to whether or it's needed for the RSS.

DUANE WESSELS: So, Russ, I think one way I see it playing it is the work party could, for example, define an algorithm or whatever you want to call that, to –

given some RSSAC002 data, this how you can turn that into a BPQ value or values—

BRAD VERD: For the system.

DUANE WESSELS: Well, it's an algorithm, so nothing stops somebody from applying—

BRAD VERD: For the server or for the system.

DUANE WESSELS: Right. To me, it would be the same. You would add them all up.

BRAD VERD: Right.

DUANE WESSELS: In our work party document, we could probably say that the reasons for this is to get a value for the system, but really nothing would prevent someone from applying that algorithm to just one operator.

KAVEH RANJBAR: Just a comment on that. I think, yes, it is [inaudible], and maybe that's good enough: adding them up for per-operator and then have it for the

system. But to be fair, in reality, I don't think it's a simple addition because, for example, all the [fate] sharing and all of that would affect [P], I guess—

PAUL VIXIE: [inaudible] will never work. That's what I was saying. It's laughable.

KAVEH RANJBAR: But I think that that still—

DUANE WESSELS: So then we need proposals for something better.

KAVEH RANJBAR: I just wanted to add, yes, I agree. It would be laughable, but still, I think that's a valuable and that's a reliable source for comparison. So if you want to compare, then that's good enough. But if you want to defend that as a number, like we can do that many bits or packets or queries and some sum this up, I don't think that would be defensible. But as a comparison point, I think that's more than good [here.]

DUANE WESSELS: But what we talked about was doing a calculation of the current load, the load on the system.

KAVEH RANJBAR: Even the current load, if you're at ...

DUANE WESSELS: The load on this day was this many DNS queries, this many bytes, this many bits, this many packets. I don't know what more you can do—

KAVEH RANJBAR: You're right. For capacity, that would be an issue, but for current load, you're right.

RUSS MUNDY: Paul, go ahead.

PAUL HOFFMAN: I guess I've been mostly today trying to wear my work party hat and looking at the statement of work. The statement of work with respect to BPQ ties directly to RSSAC037. I agree with what Kaveh said before, which is that the text in 037, which is what's being shared now, is all about capacity and not about load.

So, yes, we can loads together. We can say, "Here's typical loads," and such like that. That is a measurement. Somebody might use it, but I don't think that that answers RSSAC037. So I don't know if you folks now, who seem to be saying, "Great. We're done, at least temporarily, on BPQ," have thought about how the work party would even think about capacity.

What Paul Vixie said a little bit ago I thought was incredibly salient, which is that those three measurements aren't going to all be

approximately the same at any given moment under a current set of loads. I've got 10 B's, 20 P's, and 30 Q's right now. But I'm really hitting the top of my capacity on Q but not the others. I would rather that didn't come to the work party without better definition of what you want the work party to do with this or whether you even want the work party to think about capacity. Clearly, the work party could think about loads, but that's not what's in the current statement of work. Thanks.

RUSS MUNDY:

So it sounds like – now, when I read the SOW, Paul, what I am seeing is it says [inaudible] that BPQ measurement first proposed in 037. So then, looking at what's in 037, I agree that it really describes a capacity type of approach. But we can also ask for guidance as a work party from the RSSAC shepherds and RSSAC chairs if in fact they want the work party to focus on capacity and trying to develop something with capacity or if, at this point in time, doing a load-based answer to the BPQ question is sufficient for now.

So let me pose that as the work party Co-Chair to first our shepherds from RSSAC. Brad, you're one of them.

BRAD VERD:

I'm sorry to go back to this. I'm reading over 037 here and I feel like maybe we're getting wrapped around the axle on certain words. So BPQ was a methodology put in place to give a number of how much the system would cost to the Board. We all agree on that, right?

RUSS MUNDY: Yes.

BRAD VERD: It was not a capacity measurement of the system. It was part of a methodology that said, "Board, you determine what the capacity is and, based upon what you want the capacity to be, you use this methodology to come up with a cost." BPQ were three of the four factors, the fourth factor being the level of risk that the Board was willing to accept.

So I feel like – and maybe I'm wrong – that, if BPQ – again, I [inaudible] for myself; I'm surprised I'm defending this so much, but I'm trying to be devil's advocate and be neutral here – is what we use, if you take the daily load, not the capacity – don't need to know the capacity of everybody's system – and figure out the cost of that load, as it reads here "using industry-standard cost determination methodologies, third parties should be able to estimate the cost of a running system for any desired service capacity values of B, P, and Q" – whether or not that sentence is true, I don't know. But if it's true, you should be able to measure the cost of your daily load. Do we agree on that? I see heads nodding but I don't hear yeses.

UNIDENTIFIED MALE: Yeah.

BRAD VERD: Okay. So, if that's true, then we should be able to figure out what the daily load is – the daily load, the average load. For every day of the past

two years, we can say, "This is what basically the rough cost is," and then say, "Board, if your risk level is 100 times X, here's the dollar figure." That's kind of what this says, right?

UNIDENTIFIED MALE: [inaudible] you're getting, yeah.

BRAD VERD: That's what this says, right?

UNIDENTIFIED MALE: I think so.

BRAD VERD: Okay. But I her BPQ is a determination of capacity. I don't think it's a determination of capacity. The capacity is dependent upon the risk that the Board is willing to take. It is a factor in the equation that leads to the capacity determination.

Am I – no? Am I wrong? Tell me I'm wrong, Paul.

PAUL VIXIE: I was totally with you until that last sentence.

BRAD VERD: Okay.

PAUL VIXIE: So it is meant to be a strict technical capacity measure because those will be the three variables which, once you boil down whatever traffic you chose to throw at the thing, it's going to be one of those three thing that you have to exceed if you want it to stop answering.

BRAD VERD: Let me rephrase what I said. BPQ was not meant to be used in this room to determine the capacity of the root server system.

PAUL VIXIE: No.

BRAD VERD: That's what I meant to say.

PAUL VIXIE: As—

BRAD VERD: [inaudible] the desired capacity because desired capacity is what the Board wants to know what the cost is. We wanted to avoid giving them a dollar figure because it said it really depends on what the level of risk you're willing to take is.

PAUL VIXIE:

So there are a few problems that are not amenable to massive over-commit, if you can afford it. And we have relied on massive over-commit to keep the root server system looking as good as it has for as long as it has. We didn't do any kind of careful engineering study about exactly how thick the wires needed to be or how much this or that. We just got as much as we could and hoped for the best. Clearly, they would like us to stop doing that, and I don't blame them a bit.

But I did like what you said about using the measured load of the system and perhaps the measured capacity of the busier nodes on the network as a way to somehow estimate that, for this amount of money – this is how we're spending now – this is the amount of over-commit we effectively have in practical terms, not the technical “add them all together and divide by the number of things” or “find the finished bottleneck” or any of that. But if we could use measured capacity of the average node or of all nodes in some way and then use the measured load on the whole system to inform the Board that, for this amount of money, they can get this amount of safety or this amount of headroom, and then, from that, come up with a scale that says, if what you want is more than that, then here's the price sheet on what it would cost, then I would say that all of this had been worthwhile.

But I also, just before you call the next break, want to break a rule. You said that you can't mix baseball with root name service, and I do want to say that BPQ got us out of the inning. And that was all it was supposed to do.

RUSS MUNDY: Wes, I think?

WES HARDAKER: Am I on deck?

BRAD VERD: [inaudible] was the seventh inning stretch?

WES HARDAKER: Well, at the risk of us sending us into a tiebreaker, because I think – I'm going to clear the chair so I can duck underneath the table here in a second. We have this horrible assumption that we keep making – we've been doing it for a decade and it probably should stop – in that we believe that the risk factor is directly associated with the load of the current system, when the reality is that the attacks have proven that it's a totally unrelated metric. So you can't take the current load, multiply it by ten, and think you're safe. They're totally unrelated at this point.

BRAD VERD: Look, I agree, and it's not black and white. I think all three of those of BPQ – bandwidth, packets per second, and queries per second – are all different resources, and highly different dollar amounts between actual bits over the wire, buying bandwidth, to routers that can handle packets per second, to queries that my server can handle. All of those are addressed differently.

But, again, if we try to engineer a solution that gives you a black-and-white answer, I don't think we'll ever get there. You got to find something that is representative of the system. I think that was the goal here.

WES HARDAKER:

Right. So my point is not that that's not true. There is a cost that we can calculate for running the system not under attack. That's easy. I think that's been done in the last hour. We're done.

To state that you can calculate the risk factor when people are sending bad queries, whether they're bad queries or straight bandwidth or anything else, is totally unrelated to the current load.

KAVEH RANJBAR:

So I don't disagree, but just to recap how we got there, it wasn't started only from the financials. The actual previous question, before we got to how much it's going to cost was, if you remember, one of the things that Cherine came up with when he became the Chair. One of his early questions was, "What if the root server goes on someone – let's say Congress – calls me and says, "Hey, there's a hearing. Why didn't you provision? You said you're in charge of this thing. And we are not. That's a different story. Another hearing. Come and explain to us what did you do as people who got money for this and were responsible."

Basically, today they don't have any answer. With this, basically what they would say in a congressional hearing or any other body who might ask them is, "[inaudible] we asked a group of expert – the people who

run this for 30 years – and they told us a distributed system is too hard to measure. That’s the best way to measure it. We got the numbers. We made sure we provisioned for 100 times more than what the current load is.” So that would be something that they can answer, and I think that’s a meaningful discussion, although from a real technical, it is really far. But still, that shows a chain of responsibility and how you can handle it because, at such a level, how can you really handle these things? I mean, if you have better measures, for sure we should give it to them. But I think provides a solution” BPQ.

RUSS MUNDY:

Paul?

PAUL VIXIE:

I think Kaveh is exactly right about that. This is an understandable answer to a rational fiduciary question. But I also think Wes is incredibly on point. We used to have bar discussions and then drinking games about designing the perfect, unstoppable DDOS. Then it got serious and we started doing that as tabletop exercises, some of with various members of our military industrial complex. It’s really not funny anymore, but the fact is there is no correlation between how much ICANN invests and how much safety they will get because a lot of people, not just in this room, can [inaudible] be perfect, unstoppable DDOS, no matter how you spend.

So, if those people become our adversaries, we’re going to get a different question from different people than we’re getting now. But we’re not. We’re getting a rational fiduciary that has an easily

understandable answer that more or less goes to what Kaveh said. We asked a bunch of people with a couple of decades of experience, and they said, “This is reasonable and customary if that’s what you want.” I don’t think we should tell the ICANN Board about the perfect, unstoppable because then they might think they should solve that problem.

BRAD VERD: Maybe you solve both issues with putting a start next to it and a footnote that says, “Does not apply in attack mode.”

UNIDENTIFIED MALE: [inaudible]

RUSS MUNDY: Jeff?

JEFF OSBORN: I just have to tell Paul that it’s disturbing that we’re basically being used for plausible deniability by the Board, but that’s really the real politick here. Because that’s the real politick here, I’m always trying to think in terms of, when we produce something, what is the first question we’re going to get asked? It probably isn’t, like they’ve said, can you withstand five times your current load? But somebody on the Board wants to ask this question [inaudible]. “Can you handle a sustained two-terabit attack for an hour—

UNIDENTIFIED MALE: [inaudible]

JEFF OSBORN: Right. So I think we want to have something that at least comes to addressing that, rather than, “Can you handle three times your current load?”

What those levels are and how you determine them and how you split them across catchments and all the rest of that is really complicated, but we need a bite-size answer to give them that is credible.

FRED BAKER: Point all of that at any one node, and the rest of them will be just fine.

JEFF OSBORN: Or distribute it evenly across all of them, and all of them will be fine. But we need to come with one of those answers.

RUSS MUNDY: I’d like to go back to Paul Hoffman, who raised the question a little bit ago with respect to capacity and the wording in the SOW and the wording in RSSAC037 and see if you’re in agreement, Paul, that doing this on an RSSAC002 stats basis – are you satisfied with that answer now? Or do you still think we have an issue?

PAUL HOFFMAN: So I'm fine with doing it as long as we label that as load and as long as we decouple it from RSSAC037, unless you folks decide that 037 really means load. So either RSSAC says Section 5.5.3 is about load and then that's fine. Or the statement of work simply says "Let's look at BPQ from load and ignore capacity. We (RSSAC) will deal with capacity later." Either of those works just fine.

DUANE WESSELS: I guess a point/question for the room. I think we have [inaudible] examples in the past where work parties have struggled a little bit to answer the things that were in their statements of work. But they did the best they could, and they wrote the report, and it was fine. I feel like we don't probably need to talk about amending or modifying the statement of work as it is. We can proceed with this line of work in the work party and deliver the report to RSSAC and be in good shape.

Does everyone agree with that? Or would it be better to somehow formally amend the statement of work?

RUSS MUNDY: The statement of work for this particular bullet says, "Refine the bandwidth, packets per second, queries per second methodology first proposed in 037." As we refine it, I think what we do in our report is describe it as a load instead of the inferred capacity type of thing that RSSAC037 said.

Does that sound reasonable to anybody? Yeah, Paul? Vixie.

PAUL VIXIE:

I wish you hadn't said that. So I think that you should do that, but I think you should have a footnote to say that, now that DNS-over-TLS has been standardized and we're going to start to expect to see a lot of that kind of flow between the recursives who are our clients and our servers, there will come a time when our capacity, in terms of how many TCP connections we can start per second, how many we can close per second, how many we can maintain overall, and whether we respond correctly to a condition where we're out of resources and we've got to start closing them – all of that will have to be studied in a follow-up.

But other than footnoting it, I wouldn't want to see [you] try to do that right now.

RUSS MUNDY:

Okay. Thanks, Paul. I think this also gets to some of the discussions we had in the first session on this, and that is the values that get put in to the various metrics that get written down. So far, we're planning on the work party to be working on that later on. I think that would probably be the correct time to make sure that the footnotes get included properly, too.

Does that sound reasonable?

Okay. I think we may have achieved our objective of getting clarification on how the work party should be handling BPQ. I'm feeling dramatically better-guided than before we started this session because doing it as we've described I think is probably a handleable thing for the work party to do.

I don't have anything else that I wanted to raise with BPQ. Duane, what about you?

DUANE WESSELS: No.

RUSS MUNDY: No? Okay. So do others have thoughts that they want to put in? We have about 20 minutes before the official break. So we're actually a little bit ahead of schedule.

If not—

BRAD VERD: Do we want to talk about latency or [inaudible]?

RUSS MUNDY: I think we probably could start that.

BRAD VERD: If it isn't out of scope, I'm just curious. Do you have an estimated end date when this – I'm just curious. In your head, just even vaguely. Is it summer? End of the year? End of the decade?

DUANE WESSELS: In my head, I don't. This meeting was the deadline that we're all working to right now. So past this, I don't remember any deadlines, to be honest.

RUSS MUNDY: Yeah. And I think we all consider it a very, very important, very high-priority task. But I think that, given that we've hit the beginning set of things that will be defined as metrics, that's a great first step. When we get to defining the values that go in there and how the measurement will occur to determine whether or not the metrics are being met, that's going to be a lot harder than what we've had to deal with so far. Personal opinion.

DUANE WESSELS: So the statement of work says August 1 is our deliverable deadline.

RUSS MUNDY: Oh.

DUANE WESSELS: That's what it says. I don't know if that—

BRAD VERD: [inaudible] follow on from this. That August ...

RUSS MUNDY: Oh, thank you. Thank you. Starting our discussion on latency this morning, do we want to try to continue a little bit more down that path? I forget what our sequence for tomorrow was.

BRAD VERD: [inaudible]

RUSS MUNDY: Let's see. Where was it?

DUANE WESSELS: [inaudible]. I think the first one was actually available, but we can go out of order.

RUSS MUNDY: Sure.

BRAD VERD: [inaudible]

DUANE WESSELS: [inaudible] latency, actually. It is latency.

RUSS MUNDY: Oh, it is latency first. Ah.

DUANE WESSELS: Would you like me to ...

RUSS MUNDY: Go ahead.

Okay. 15 minutes on latency, and then break time. How's that?

UNIDENTIFIED MALE: That works.

DUANE WESSELS: Okay. So I'm not sure who's running the screen. Steve [inaudible], can you bring up that doc again? This doc. We're going to talk about the RSO metric latency. It's probably on Page 3 of that document or so.

So we've already spent a little bit of time today talking about latency, so maybe this will be faster than otherwise. The document here I think explains where this came from because it's derived from the statement of [work] [inaudible] to assess the performance that each RSO provides. So we take latency as a pretty simple metric.

It has here listed what I call sub-metrics. So that is, for each root server operator, you could imagine doing these latency measurements over this different transport permutations – v4 UDP, v4 TCP, v6 UDP, and v6 TPC. It doesn't have to be that way. It's just a proposal that that's how it's done in other contexts. Maybe we want to do something similar here.

I see Liman has a question.

LARS-JOHAN LIMAN: Yes. With the experience of trying to form measurements according to the Applicant Guidebook, doing measurement over TCP is kind of interesting. So be careful that you actually run tests before you specify in the final document how to measure thing over TCP. Just make sure that you measured this doing it this way, blah blah, blah, and that you have tested to do that.

DUANE WESSELS: Okay. Yeah, I think that, as we talked about earlier, we could take for granted at this point that, if there's a measurement platform, it will be do things like test its basic connectivity so that it knows when it's on the Internet or not on the Internet. You could imagine also, if all probes or some number of probes above some threshold – maybe above 50% -- all have timeouts on a particular TCP metric, they all get thrown out or something like that. So you do some sort of smart things to try to detect these types of [inaudible] failures or difficulties.

Is that where you were going with this?

LARS-JOHAN LIMAN: Not really. The thing I'm comparing to here is that the Applicant Guidebook has requirements for the TCP capacity of the service. To measure the TCP capacity on the server is not easy because there's so many parameters involved, and loading a server with UDP traffic or loading with TCP traffic are two totally different things. The guidebook gives you now guidance as to how to measure this. I'm quite certain

that they haven't thought it through. They just rubberstamped UDP to TCP without thinking through how to do that.

So the only thing I require is that, when you measure things here, you be sure that you have tested that you actually get the values out of this that you think you're getting out of it by running code. If you do that, I'm fine. Outliers and stuff are not a problem. And connectivity is not the problem. The problem is actually the interaction between the client [of] the server and what do you measure. Do you measure from sending the TCP message with the query to receiving the TCP message, or all the segments [for the TCP message] for the answer, or tearing down the session, or .. Just be sure that you have the details in there.

DUANE WESSELS:

Okay. So, yeah, it sounds like you're advocating for a very detailed description of how to do a measurement and then actually to test your measurements before putting them into production and maybe even finalizing the document and that sort of thing.

LARS-JOHAN LIMAN:

[Yeah.]

DUANE WESSELS:

Okay. Andrew, was there a question from someone online, or a hand?

[RUSS MUNDY]:

Paul, did you have your hand up?

PAUL HOFFMAN: I did but I lowered it just because Liman finally hit the right thing, which is that these are not to be measurements. The Applicant Guidebook measures capacity – I'm sorry. Pretends to measure capacity, which is insane. And these measurements are supposed to be latency in the real world, not capacity.

LARS-JOHAN LIMAN: Yes, but there are still so many moving parts in TCP that you really have to specify under which conditions you measure this.

PAUL HOFFMAN: Fully agree.

DUANE WESSELS: Okay. Thanks. So there's some text here but it's very terse. It doesn't go into a lot of the details that we've just been talking about. So [inaudible], but the idea is that, at some interval – for example, maybe every minute – the probe clients send a DNS query for each of the sub-metrics above and it measures the time between sending the query and receiving the response and records those.

KAVEH RANJBAR: Since Paul is online, shouldn't we add latency over [quick] with the [inaudible]?

DUANE WESSELS: Somebody cut his microphone off, please.

PAUL HOFFMAN: No!

DUANE WESSELS: Well—

UNIDENTIFIED MALE: [inaudible]

DUANE WESSELS: Right. I think that's an open question: the extent to which this work party document should say anything about future transports. Maybe it wants to leave that for a future work [party], or, "I [acknowledge] that future transports will exist and will have to be addressed," or ... I don't know.

UNIDENTIFIED MALE: Paul, go ahead.

PAUL HOFFMAN: What Duane just said, which is we should be measuring the current DNS. We actually have no idea when DOT between recursive and authoritatives will actually be considered seriously. It has come up in the IETF fairly often, and it never moves forward. I think that, at most, a

footnote saying, “Future transports will be considered in the future,” would be completely sufficient.

DUANE WESSELS:

Okay. All right. So we talked about the first paragraph, which talks about how often we might want to make these measurements. There’s a blank here for dealing with timeouts. There would be some TBD timeout value – for example, maybe four seconds – and, if it exceeded by that much, then it would be considered a timeout, not an [inaudible] measurement.

Then it talks a little bit about how you would take all the measurement from all the probes over some other time period – maybe one day or maybe one hour or whatever – aggregate them, and calculate a median value or perhaps other percentile values, and [report] on those.

So, if you imagine we’ve got 13 server identities for sub-metric transports and tens or hundreds of probe servers, you’re talking about quite a large number of measurements to aggregate and calculate a median over.

BRAD VERD:

Is there any discussion on the geographic diversity of the probes?

DUANE WESSELS:

So we haven’t had the discussion yet. It’s acknowledged in this document that it’s something that needs to be addressed, but we don’t have anything proposed at this point.

BRAD VERD: It needs to be addressed.

UNIDENTIFIED MALE: Yeah.

BRAD VERD: Okay.

DUANE WESSELS: It says something like you might imagine not having more than one probe in the same AS number or distributed among continents, but there's not a lot of specificity there yet.

BRAD VERD: There's got to be some – I don't know – industry best practice that represents the client base because the client bases or not all in one spot. You're going to have different regions throughout the world, and maybe it's just a couple of probes in each of those regions and you're done. But you've got to identify what those regions are. Even then, it's not perfect because of the network topology. I get all that. But it seems like you want best effort there somehow.

DUANE WESSELS: I think [GTI] Guidebook as some language about this, too, talking about placing probes in regions with lots of actual clients and things like that.

Okay. So, for that sort of strawman metric, here's the list of things that are yet to be decided. You've got the measurement interval, specifics about the DNS query, the name type, etc., and other parameters, timeouts, do you want to have retries if the query is timed out, the aggregation interval, percentiles, and then, lastly, which we've – good thing it's last because we've kind of agreed that it should be the last thing to discuss: the thresholds for what good looks like.

So shoot more holes in this proposal if you feel like [it].

PAUL VIXIE:

When we talked about latency this morning, I mentioned that there are several components of latency as measured by a specific test. One will be just transit delay, just how long it takes – speed of electrons, speed of photons. There's a minimum amount of time that a ping would have taken on an unload network.

Then there is this additional latency, which is various queues – port input queues or socket buffers or in various places where data is held while waiting for its chance to get physical access to some media somewhere. To the extent that you can use deviation in the results you're seeing to be able to estimate what is the floor, what is the actual speed of propagation latency versus the variable component, which represents congestion, I think they're differently interesting because, if for, let's say, all of Europe a certain root nameserver has a minimum of 100 milliseconds, that means you're crossing an ocean to get to that root server if you're in Europe. That would tell us something.

If, on the other hand, the minimum that we're observing is 25/30 milliseconds – so it's in metro – but we're seeing huge swings of an additional 50 to 100 milliseconds, that would tell us that somebody either has buffers that are too big or traffic that is too much. And that tells us something differently interesting than the floor.

So, to the extent that you can work into these metrics a way to estimate which part is due to physics and which part is due to engineering or economics, I would love to have those separately thresholded. Thank you.

FRED BAKER: A third possible explanation would be a multi-path route.

UNIDENTIFIED MALE: [inaudible]

PAUL VIXIE: Yes. A multi-path route would exhibit similar characteristics, but I think they can be differentiated. I think you'll be able to tell which one is from which cause. There's a whole lot of way smarter measurement people, including Duane, than me who would know the exact references to [inaudible] in order to make sure you were getting that part correct.

DUANE WESSELS: So thanks for the question, Paul. I think one way you could address that is as suggested here. One of the values that could be reported would be,

say, the 5th percentile, over the course of the [day]. That could be your minimum. Or you could take 1st percentile or whatever. And you could a median, and you could have different what-good-looks-like for those two different aggregates. It seems that would address your concern. Yeah?

TBD?

PAUL VIXIE: It would be a start. It's certainly worth proceeding on that basis to see if a better idea occurs to anybody when they see the next draft.

DUANE WESSELS: Another thing that would be [inaudible] – I'm reluctant to sign the work party up for this, but one thing that somebody could do if they had a lot of time and resources is you could take measurements from an existing platform, like RIPE Atlas, and you could do these sorts of aggregations and see what comes out and see what it looks like. Might be interesting.

BRAD VERD: Just a thought. This – I don't know. Maybe I'll move the chairs to climb underneath the table, Wes. But maybe [inaudible] what good looks like is different on a per-region basis.

DUANE WESSELS: Per region of the probe client?

BRAD VERD: Yeah. I don't know. You start thinking about where the infrastructure is and where it's not. Good in where it's not might be completely different than what good looks like from where the infrastructure is.

DUANE WESSELS: Certainly you do that. Probably the first question I would have is – go ahead.

BRAD VERD: It just occurred to me as you were saying this – taking the 5th percentile or whatnot – that I think maybe if you define the measurements, run the measurements, for a period time and create a histogram, you start seeing what good looks like. Maybe we have an idea of what we think good is, and then you go in and you run it and report on it and you look at it and you're like, "Wait a minute. We were way off," or, "We were way too high. Good is way down here, but in this region, good is way up here." That's just a thought.

DUANE WESSELS: So if we go to the work of defining the metric, we can start collecting that data and look at it without having defined the thresholds first.

BRAD VERD: I'm [inaudible]

DUANE WESSELS: Yeah.

WES HARDAKER: Did you just agree with me four hours ago?

BRAD VERD: For crawling under the table.

DUANE WESSELS: It should be obvious, I guess. But what Brad said – in addition to be it being higher in some place, it's going to be a lot lower in some places. Showing that you failed because of a probe in Namibia hitting something in Lesotho? There's all kinds of slack in [an earth station].

Hiro?

HIRO HOTTA: Sorry for a [inaudible] question. Can latency be a metric for the whole root service? If so, will it be measured?

DUANE WESSELS: So you're jumping about half-a-day a head of us. But it's a good discussion to have now. One of the things that the work party is not decided on or hasn't made progress on is how do you measure the root server system versus how do you measure an individual operator?

So, for example, one thing you could do is you could do all the individual measurements and then collect them somewhere and then use math to come with a value for the system as a whole. You could do, again, 5th

percentiles, medians – whatever. Averages. That’s one way to do it. Or maybe you want to have them as totally separate measurements, maybe even done via different techniques.

So we don’t have anything to propose at this point. We don’t have any – well, I guess if we ... We’re going to go through this tomorrow afternoon, but maybe this document probably leans a little bit towards the first approach, which is where you do measurements to all the operators individually and then have an algorithm which takes those and turns it into a single system-wide measurement.

But I think we’re totally open to other ideas. For example, one of the ideas that I’ve put forth a couple times is you could install off-the-shelf recursive nameserver software in these probes. You ask your question of the recursive software and let the recursive software do its thing, where it chooses which authoritative nameserver to query and it does its thing of retrying and choosing the transport, choosing v4/v6, and that sort of thing. So that’s another way you could imagine this working. That has some advantages, but it has some disadvantages. So I think we’re open to either one.

RUSS MUNDY:

Just a quick comment here that, in terms of, again, the measuring mechanism, I think, at least for Duane and myself – and I think the work party has generally supported this when we talked about it – whatever ends up being the measurement mechanisms, we would very, very much like it to be based on something that is really available, e.g. open source, that other people could pick up and make use of, regardless of

whether or not the official measurements are done by a particular group or a particular set of probes. But to make it as open and available and as freely available to anyone that wanted to do this type of measurement is an objective where we want to get to.

DUANE WESSELS: I think we're at breaktime, probably. Good time to stop?

RUSS MUNDY: I think so.

DUANE WESSELS: Yeah?

UNIDENTIFIED MALE: Yeah.

DUANE WESSELS: Okay. Thanks.
So back here at 4:30?

BRAD VERD: Yes. Back here at 4:30 for a quick wrap-up.

FRED BAKER: Wrap-up/review of Day 1 ... oh my goodness. Anyhow, let me just run down the agenda for today. We had a working breakfast. We actually ate. This is good. We talked about RSO independence and, Ken and Andrew, you've got a [project you're] working on. I don't know whether we'll actually [vote on] that on Thursday or not. We'll see how much it is and whether we are comfortable with that. We may wind up doing an electronic ballot. That would be next week, right, Carlos?

CARLOS: [Yeah].

UNIDENTIFIED MALE: [inaudible]

FRED BAKER: Yeah. So something like that. We talked about RSO funding, and that is a project that is just kind of getting started. But you've got some homework to do.

JEFF OSBORN: I'd say everyone has some homework to do.

FRED BAKER: Okay. Well—

JEFF OSBORN: By Marrakech, we need to have feedback on what [inaudible]

FRED BAKER: Microphone, please.

JEFF OSBORN: I'm sorry. Then Ken came up with what is a really useful term we're going to want to use, where those who are not necessarily getting money could talk about an MOU rather than and SLA. Is this reasonable to talk about in public, or was for – okay. There's some clever thoughts, like, if you had to do an SLA from a government agency, it's possible you'd have to go so many levels up at such a cost that the money was wildly outweighed by the cost of just going through your organization, where an MOU may be able to be signed at a much more local level.

So it's intriguing if nothing else.

FRED BAKER: [inaudible]

JEFF OSBORN: Correct.

FRED BAKER: [inaudible] MOU [inaudible]

JEFF OSBORN: Paul signed an MOU years ago – exactly – that started—

FRED BAKER: [inaudible] somebody else [inaudible]. And I and [inaudible] have [inaudible] right?

JEFF OSBORN: Mm-hmm.

So anyway, I was tasked with presenting something and did my hour-and-a-half. Am I being voluntold I'm—

UNIDENTIFIED MALE: [inaudible]

JEFF OSBORN: We need to do that, and then there's a homework bit. I'm just—

UNIDENTIFIED MALE: [inaudible]

JEFF OSBORN: Should I volunteer for being – when this thing lands, I'll catch it?

Okay.

FRED BAKER: I think you are probably the best qualified person in the room to catch it.

JEFF OSBORN: Flattery will get you everywhere.

FRED BAKER: Hey, I pay you, too.

JEFF OSBORN: Cash helps.

FRED BAKER: Okay.

WES HARDAKER: Jeff, talk to me at dinner, would you?

FRED BAKER: We talked about the Empowered Community, and, Brad, you made some comments to the effect that that might be of value for us to consider for the RSS [NG]—

BRAD VERD: Yeah. It was just to keep it in your mind as we start working through some of this implementation 037 and everything else to see what it is that we want as a group. I'm not sure what that answer is because I hear concerns. I see some things that can mitigate some of those

concerns, so I want to make sure that those are on the top of everybody's mind.

FRED BAKER: Okay. And, Duane and Russ, we spent a fair amount of time this afternoon with the metrics discussion. My sense was that you achieved what you wanted to achieve. Is that correct?

RUSS MUNDY: Yes.

FRED BAKER: Okay. And we're having a wrap-up. So just going through the agenda, I think we've actually had a reasonably successful day. I want to thank you guys for your involvement.

Tell me about dinner. Carlos?

CARLOS: Dinner tonight is at 7:00 P.M. We're going to Uncle Julio's, which is a Mexican restaurant here in Reston Town Center. Let's meet in the lobby at 6:45 and then we can walk over. We have a set menu for everyone. There are multiple options and an open bar. So be judicious about that.

UNIDENTIFIED MALE: Did you put camarones brochette on the fixed menu?

Oh, man!

FRED BAKER: Okay. So, with that, I think we're done for today.

BRAD VERD: Andrew, Carlos, or Steve – somebody – were there any other action items that we missed during discussion?

UNIDENTIFIED MALE: No, I think you covered it.

BRAD VERD: Okay.

FRED BAKER: Okay. Cool. I'm going away.

BRAD VERD: Adjourned. See you at dinner.

[END OF TRANSCRIPTION]