KOBE – RSSAC Work Session (4 of 8)
Sunday, March 10, 2019 – 09:00 to 10:15 JST
ICANN64 | Kobe, Japan

FRED BAKER: Okay. I think it's 9:00, so let's come to order. We've got a bit of a change to the schedule that we want to make. We don't think we're completely done with the concept paper. Is that a fair statement?

So will have the RSS metrics discussion during the first session, and we're about to have that, and then the second session which is on the schedule dedicated to RSS metrics. We're actually going to go back to the concept paper discussion.

Does anybody have heartburn with that?

Okay. So, Steve, can I hand the ball to you?

STEVE SHENG: Thank you, Fred. Good morning. This is Steve Sheng, ICANN staff. This is the second meeting for the RSS Metrics Work Party since the first meeting. We went over the scope and discussed approach a bit. The key element of that is selecting a co-chair.

At the time, Duane and Russ raised their hands to express interest. Also, Daniel expressed interest to be a backup should one of them not want to do it. But we've confirmed the interest, so I've sent out an e-mail earlier this morning [about how] RSSAC appointed Duane and Russ as the work party Co-Chairs.

So, with that bit of housekeeping out of the way, I think anyone here who is not – oh, so we have Paul here, a Caucus member, and we have a few people on the line. So please state your name before you speak. For remote participants, they can speak as well.

Okay—

UNIDENTIFIED MALE: No. They can just write their questions in the chat.

STEVE SHENG: Okay. But does their voice [channel] through the system?

UNIDENTIFIED MALE: I'm afraid [inaudible].

STEVE SHENG: Oh, okay. All right. So it's a listen mode. Okay. So this is for remote participants: if you have any questions, please type in the chat, and [you'll] get followed along the way.

With that, let me pass this over to Duane.

DUANE WESSELS: Okay. Thanks, Steve. So Russ and are happy to be your Co-Chairs. I'm going to do most of the talking today. Russ is just off the plane and just joining us. So the plan today is for me to chair the meeting.

This is a Caucus meeting, so if you're in the room here and if you're a Caucus member, you're welcome to come sit at the table and use the microphones to ask questions.

So we'll start with an overview of the agenda. Are we having Adobe Connect problems, or … oh, there it is. Okay. So, briefly, the agenda for today is to review and discuss the work approach. We'll talk about the statement of work. We want to spent probably the most of the time going through reviewing some of the existing documents and materials that are related to this work party, and then we'll discuss the next steps, including our schedule and our first milestones for this work party just in advance of the RSSAC workshop in April.

All right. Next slide. Got to make that bigger. This text is from the statement of work. So the task before us today is to define the system-wide, externally verifiable metrics that demonstrate the root server as a whole is online, serving correct content and timely responses to end users, and define measurements to ensure the root server operators are meeting a minimum level of performance. I won't read the rest to you.

Yeah, you can go to the next one, Steve. Thanks. Also, the statement of work talks about this idea that's in the RSSAC037 document, which is to refine the bandwidth, packets per second, and queries per second measurement methodology.

Then what the statement of work says is, depending on the results of the work, it may result in updating RSSAC001 or it may be something entirely new. That remains to be seen. I think we'll talk a little bit more about that today.

We also have RFC 7720, which is a companion document to RSSAC001. It's even possible that the output of this work party would result in an [Internet] draft or RFC as well.

So the Metrics Work Party had a meeting a couple weeks ago, and one of the things we talked about there is we wanted to make sure that all the work party members would review the statement of work and understand the scope. The idea we had was to have everyone take the statement of work, the items that are listed there, and write their interpretation of what those items mean.

So I believe Steve or the staff is going to put together a Google sheet or document where the work party members can go in and fill in their answers, if you will. At our next meeting, I would like to review those and find places where people had different interpretations of some of these items and talk through those so that we're all on the same page.

In addition, what we want to do as a first goal is to brainstorm the list of metrics that should be included in this work party. We will avoid, to the best of our ability, talking about any limits or particular [values] or thresholds. We just want to talk about what the metrics should be, not what the limits on those should be.

We expect this is going to be an iterative process. We'll have numerous meetings and feedback from RSSAC. So these metrics will be refined over time until we get to our ending point.

All right. Next. [inaudible] So this is what I think we'll probably spend most of our time on today. This list here is five or so documents where RSSAC and other places currently define metrics or expectations type of things for either root server operators, root service, or for TLDs. We have RSSAC001, which I think is titled Service Expectations of Root Server Operators. That has RFC 7720 as a companion document.

We have RSSAC002, which describes measurements that the root server operators are expected to perform and publish. We have RSSAC024, which is some work that was done prior to the RSSAC037 governance model. RSSAC024 describes, I think, technical metrics of potential future root operators or something like that. So I think there's maybe some idea we can pull from that.

RFC 7720 we already talked about. Lastly, the Applicant Guidebook has some technical specifications for TLD operations, which may be helpful for this work party as well.

So that's the plan. I'll ask Steve to put up those documents one at a time. I just to go through those items and just remind people of what's in those documents and get people thinking about whether these sorts of things are within scope of the current work party or not, thinking about if these documents need to be updated by this working party or not and that sort of thing.

Okay. All right. So first on the screen here is RSSAC001 version 1. So that means that this document has not been updated since its initial publications. It lists these expectations of the root server operators.

My idea was just to focus – I think they're succinctly captured in the recommendations in each section, so maybe, Steve, you can scroll until we find some bolded recommendation text that we can read.

So here's—

STEVE SHENG: Let me go towards the end.

DUANE WESSELS: Are they all listed at the end.

STEVE SHENG: Yeah.

DUANE WESSELS: Yeah, I think so. Yeah. Okay. For example, the first one says that RSOs are to publish operationally relevant details of their infrastructure, including locations and addressing, and routing information – autonomous system information.

The next one says that RSOs will deliver the service in conformance to IETF standards and requirements required in the companion RFC and any other standards as appropriate.

The next expectation is that RSOs will adopt or continue to implement the current DNS protocol and associated best practices through appropriate software and infrastructure choices.

3.2. Individual root servers will serve accurate and current revisions of the root zone. Individual root servers will continue to provide loosely coherent service across their infrastructure.

I believe what that means is best effort in serving the current version of the root zone so that all sites have the same data, more or less.

[3.2D]. All root servers will continue to serve precise accurate zones as distributed by the root zone maintainer. Individual root servers are to deployed such that planned maintenance on individual infrastructure elements in possible without loss of service availability.

Infrastructure used to deploy individual root servers is to be significantly redundant such that unplanned failures does not cause the service to become generally unavailable to the Internet.

Each operator shall publish documentation that described the commitment to service availability through maintenance scheduling and notification. Operators will make all reasonable efforts to ensure that sufficient capacity exists in their deployed infrastructure to allow for substantial flash crowds or denial of service attacks. Each operator shall publish documentation on the capacity of their infrastructure, including details of steady-state load and the maximum estimated capacity available.

Operators will adopt or continue to follow best practices with regard to operational security.

Well, there's more of these than I expected.

Operators shall publish high-level business continuity plans with respect to their root server's infrastructure.

So that's disaster recovery stuff.

Operators shall publish document that describes key implementation choices, such as the type of DNS software used in the interest of diversity of implementation choices across the system as a whole.

Each operator will adopt or follow best practices with respect to monitoring elements within their infrastructure. The operator will continue to perform measurements of query traffic received and shall publish statistics based on those measurements.

So that sounds like RSSAC002 to me.

Individual operators will continue to maintain functional communication channels to coordinate and so on. Communication channels are to be tested regularly. Individual operators shall publish administrative and operational contact information to allow others and interested parties to escalate service concerns.

That's it. So as I read through those, some of those I think talk about things that are measurable, and some are maybe not so measurable. I think some of those things can be put into the Metrics Work Party but others not so much.

Anyone else have comments or thoughts on the expectations that we already have in RSSAC001 and their relevance to this work party?

FRED BAKER: Wes?

WES HARDAKER: I just wanted to say this is an open Caucus meeting or open meeting, so please, if you're in the back of the room and you have comments about this topic, this is not [open] just to RSSAC members, unlike some of the ones yesterday. This is open for others to comment on.

[PAUL HOFFMAN]: So I have a question on observable versus verifiable because you had said earlier one of the things that we're working on is to have the metrics be verifiable. So some of these things are observable but not verifiable.

How do you feel – or should we wait until we do the survey to see if people are interested in that? Because, if we limit ourselves to verifiable – we have everything that's here, everything that is observable, and then everything that's verifiable. I don't know if we want to cut down that far or not.

DUANE WESSELS: Can you give an example or give more detail of what the difference is? Are you thinking something algorithmically verifiable versus just observable by somebody looking at a web page, or …

[PAUL HOFFMAN]: Right. So it is observable that somebody has said what they're doing. It's not verifiable that they are doing it. I mean, you can verify that they said it. But I'm just wondering whether we're – I don't want us to get ratholed, but I also don't want us to be too broad on – because some of the things – another example would be that you can observe that this operator was under attack on a certain day, but you can't verify the level of the attack. Things like that.

DUANE WESSELS: So I guess my hope is, or the Metrics Work Party, we focus on things that are technically measurable and less on the things that are just observable. That's where my head is.

[PAUL HOFFMAN]: Right.

UNIDENTIFIED MALE: Similarly, if we took the first page seriously, it talks about developing metrics and the what constitutes failure in them. I'm assuming we're taking that second half and throwing it with Paul's verifiable, because firs the measurements and then the sentencing. [Okay].

DUANE WESSELS: Okay. No more comments before we move onto another list?

All right. Steve, do you want to put up 002? I don't remember if these are – are summarized at the end or not? I don't remember.

So this is RSSAC002. These are measurements that—

FRED BAKER: There's only three.

DUANE WESSELS: Only three recommendations? Okay. These are measurements that root server operators are expected to make and publish. This list here in the table of contents is probably a good starting point. We have the load time metric, which is the time that it takes to "load" the root zone. This was a little bit hard to define, I remember, but essentially it's the time from when a server receives a notification that there's a new zone until the time when it can serve the first query from that zone.

The zone size metric is the size of the root zone. That is something that has evolved over time and currently officially only to be reported by the root zone maintainer, I believe.

The traffic volume metric is essentially the number of queries per day over various transports. So over UDPv4, UDPv6, TCP4, TCPv6.

All of these measurements are actually on a daily basis, so they're not very granular, I guess. So we have traffic volume. Traffic sizes is essentially a distribution of the sizes of the messages received again over all the permutations of TCP, UDP, and v4/v6.

The R-code volume is the distribution of response codes from responses served by the root servers. The unique sources metric is just a count – actually, it's three counts – of the number of unique source IP addresses, both v4, v6, and I think v6 with some /64 mask or something like that.

So those are the seven or so metrics defined, and these are things that root server operators measure on their own infrastructure and then publish results of. These are not designed to be things that are measured externally. This is something that the operator measures itself from the traffic that it receives and then publishes as YAML files.

Liman?

LARS-JOHAN LIMAN: This is Lars from Netnod. Just a quick comment that you mentioned that the time is quite long. It's a full 24 hours. But the reason for that is to even out changes during the days. The reason for collecting these statistics is to be able to observe long-term trends. It was never intended to make momentary observations of the system. So there is a reason for this design.

DUANE WESSELS: Yes, absolutely. It was very intentional – the choices made here, the measurement interval.

RUSS MUNDY: Just one comment with respect to our work in this work party. We haven't really included a requirement to define the timespan over which the measurements are going to be laid out. But that's probably a reasonable thing to also try to address. Are we talking RSS metrics on a charted instantaneous basis or are we talking 24 hours or 30 days or six months [here]?

BRAD VERD: I would hope that the Metrics Work Party – I believe the intent of the statement of work is to define what good looks like and create technical accountability for the root server system and the root servers. So if you're looking at anything beyond – I don't know. If you're looking at daily metrics, weekly metrics, and monthly metrics, we're already beyond, to me, what [–] we failed. It's got to be – Is the system up? Is it not up? How is it performing?

As far as query loads, as far as the 002 metrics, I think these are good statistics and good metrics, but they're not necessarily related to the intent of this work party in my opinion and my interpretation.

WES HARDAKER: Wes Hardaker, USC ISI. I think that we have to take – I think that's a really good question, first off, Russ – it on a case-by-case basis because, in order to define expectations, there may be some that we need to measure over the course of a day, a weekend, or a month (probably not), and then other ones might be five minutes, other ones might be an hour, because how you define uptime typically: is it up right now and you fail if it's not? Is it up more than n percent and n percent requires measuring for a long period of time, for example, in order to determine that? But other things –can you handle a query rate of so many per minute? That's a different, much shorter metric because you need to catch spikes and things like that. Something I'm thinking very deeply technical.

So my answer to you is I think that it's going to be on a case-by-case basis, but that is an important thing to consider.

UNIDENTIFIED MALE: Brad makes a really good point. We might be making this too complicated. "Who's the audience?" is a really good question for me. We did something after years of mucking through what is the cash level of the small company. We came up with this really simple metric. Red is a problem. Yellow is worrisome. Green is probably okay, and blue means everything's great. You just walk by the CFO's door or look on his website and there's a blue square and you keep going with a spring in your step. That's kind of all anybody needed to know, whereas we used to spend hours saying, "Well, receivables are here and payables are there and your budget is another thing."

I'm just wondering, if our audience is like the ICANN Board – we're already way over all their heads with things like bits per second. So …

UNIDENTIFIED MALE: Well, having not thought about this for too long, really, I think one of the end audiences for this document will be a group who's tasked with implementing and making these measurements – a technical body whose job it will be to do these measurements.

UNIDENTIFIED MALE: Again, to what end?

UNIDENTIFIED MALE: To what end? Why are they doing the measurements? So that we know if the root server operators and the root server system is meeting its expectations.

UNIDENTIFIED MALE: So whether they're red or yellow or green. Okay.

FRED BAKER: I read this in part – I'm coming back to RSSAC037 – as this is sort of the PMMF: to do whatever the measurements are that are there and then, if there is an issue on the basis of those metrics, to be able to detect it and understand it and figure out what to do next. In the concept paper, the names change, but that's basically what's going on. So we're describing, if you will, the function of the PMMF.

Now, I frankly do think the delivery of these statistics and the understanding of these statistics is part and parcel of that. The question is whether the PMMF is limited to that, whether there are other things.

When we talk about weekly or monthly statistics, I would assume they can be derived from daily statistics. I'm not sure, in that context, the value of a five-minute metric or something like that.

But I think that's the answer that Jeff is looking for: to answer the questions that are raised in RSSAC007.

UNIDENTIFIED MALE: So, talking about the [inaudible], I think the [inaudible] should be some developers so that it can implement, actually, what we're providing. Maybe the output might be a color. Whatever. But it should be implementable, I guess. That's why I'm expecting from one of the documents. It should be translated as an RFC.

BRAD VERD: While I agree with you, Fred, about the PMMF, I don't want to conflate things, and that is your answer is accurate should 037 be implemented. This work right here we thought needed to be done regardless of 037 and could be done independent of 037. Should 037 be implemented, it would just plug in and be executed by the PMMF. So I don't wan to lost sight of the fact that we all agreed that, regardless of what was with 037 or when 037 happened, this was important work that needed to be defined.

FRED BAKER: And I very much agree. What I was going at what was the question, who's asking it. I think that's [there].

WES HARDAKER: Thanks. Wes Hardaker, USC ISI. I think one core audience is actually the root server operators that are implementing the service because, to a large extent, this document defines the expectations of the entire Internet upon the service that at least my team, for example, provides. I care very deeply that I am providing the service that is needed.

So we need to go through this iterative cycle on a regular basis to make sure that – because, if I'm not providing that service, I need to change. So this is one way of doing that, and it may turn out to be green, yellow, or red. But that has to be translated by numbers at some point.

So I know I will be, even though I'm involved in writing it, equally as involved in leading it and making sure that I live up to what's inside of it because that's my job.

PAUL HOFFMAN: I have a bit of a concern on timing, although I think it's a solvable concern. I was not at all active in the 002 v3 work, but I heard moans and groans over a long period of time when it was happening. A little bit of the statement of work says, well, we might actually work on another revision to 002. If that's really part of what we're supposed to be doing, that may also take a very long time and such. So I think we need to schedule which things go first or whether we even do it, a tricky thing because we haven't actually looked at the conclusions, which I had skipped to. In 002 v3 it says that RSSAC will reopen this in two years, and that was just about two years ago.

FRED BAKER: The comment in the statement of work about RSSAC002 came from me. What I had in mind in writing that was that, if we're discussing metrics, it might touch on the metrics that we. It might – question – touch on those metrics. So I just wanted to make sure that the possibility was there.

UNIDENTIFIED MALE: Yeah, I don't see this Metrics Work Party updating RSSAC 002 myself. I think that's out of scope. But I do think it might be appropriate to have a metric that checks the metrics, to have a metric that, if these are to be published, then the operators – you check that they're actually being published. I think that would appropriate. But I don't see us updating this document.

DUANE WESSELS: Okay. Should we move on to the next one? Guess we should have gone to the RFC-2, but we can do that in order. Ah, this is RFC 7720 then. So this is the companion to RSSAC001. I believe this is a relatively short RFC. It mostly concerns itself with following the other RFCs.

Can you scroll down to Section 2, Steve?

Okay. So the first section talks about high-level protocol requirements. It also references 001. So there is quite a bit of overlap here. For example, the root name service must implement core DNS RFC and its clarifications. Must support v4 – and I think that says v6 under the hidden box. Must support UDP and TCP. Must support UDP check sums. Must implement DNSSEC. And must implement EDS0.

So that's a pretty straightforward list of protocol requirements. Very straightforward, something that is measurable. So we could imagine maybe some metrics around these requirements.

The next section is very short. The root name service must answer queries from an entity; basically, any source [inaudible] conforming to RFC 1122 with a valid IP address. And must serve the unique root zone as referenced.

Those might be a little hard to measure, but we could tackle those, too.

Is there anything in – I guess the security considerations is probably just a no-op, right? Yeah. Okay.

So that's sort of short and sweet, something to keep in mind as we go through the metrics for the work party.

Anyone wants to make comments on what's referenced in 7720?

So the last one is the document RSSAC024: key technical elements of potential root operators. As I said before, this work was done in advance of the 037 work with the idea that we need to start thinking about what sort of standards we would hold potential new operators to.

These metrics are listed in Section 3. So there's six of them. I think there's actually quite a few.

[FRED BAKER]: [Yeah. More on the list].

DUANE WESSELS: Some of these, again, are more measurable than others. There's a reference to our existing documents of 001 and the RFC. So I was involved in the work party for this document, and I think one of the goals was – our idea was that there would a third party who would be evaluating a potential future root operator and this would be advice that they could follow and things that they should look for from some – I don't know – "application" or – I forget what we called it exactly.

So the first one is very broad. It's overall service design. The overall service design should be evaluated with its respect to utility and serving the root zone, provide as many details, yada, yada.

The proposal must be evaluated with respect to its approach in maximizing service availability. The design is expected to eliminate or minimize single points of failure. That echoes something from 001, as do a lot of these.

The candidate operators' service capacity must be evaluated for its ability to withstand DDoS attacks. Should be evaluated with respect to its performance characteristics, such as latency, service regions, and RSSAC metrics.

So far, to me, these are very broad and there's nothing really substantial or significant that we could attach a metric to, I don't think.

But let's see. Operational experience. It is expected that any future operator would have prior experience in operating similar services. That's what that's about.

BRAD VERD: [inaudible]

DUANE WESSELS: Oh, sorry, Brad. I wasn't looking.

BRAD VERD: Can you scroll back for to that last one you just read? The 324?

So I believe that "broad" is written here. I guess my question is, are there things that can be identified that maybe we need to create a metric on, for example, such as latency? It seems like there needs to be a metric there. Such as service regions. Maybe there's a metric there, like it needs to be in X number of regions versus one location.

Can we extrapolate from these very broad topics to something more granular or specific that drives towards the availability of the service?

DUANE WESSELS: Yeah, absolutely. Thank you. That's a good point.

All right. Let's continue down. The next one is about experience. We talked about that a little bit already. 322 is about an audit. There was an expectation that any candidate operator would be asked to provide some security audit details conducted by a third party. Those results would be kept private.

A future operator was to obtain its own AS numbers and IP addresses for operating the root server. There was an assumption that Anycast would be used. This talks a little bit about making sure that the allocated networks are reachable from most of the Internet and not, for example, listed in – I forget if we left this text in. But in this section we talked about how service addresses shouldn't be in black list and things like that. You don't want a tainted address base for your root server operator.

PTR records should exist – oh, here's the part about the black list. The address block should be evaluated with respect to the reputation. Peering data should be kept up-to-date and in routing databases. These are things that we could build metrics around, I think.

Address space should be accurately registered in RIR databases. 33[A] talks about the zone distribution architecture. Today, the way all the current root operators function is that they have their own, for the most part, internal zone distribution system. We expect any future operator to have a similar situation.

All right. Scroll down.

So this section talks about various aspects of diversity and ways that diversity is beneficial. It talks both about diversity within an existing operator and diversity among operators. So one of the first is geographic diversity, which we talk a lot about. There's another work party about this underway.

We talk about network provider diversity so that individual operators are not susceptible to the problems and sustained outages of single providers. We talk about hardware diversity to guard against zero-day vulnerabilities. Similarly, server diversity can refer to, again, hardware or even different models of hardware from the same vendor.

Operating systems diversity is something that's called out in RSSAC001. This may be an example of something, Paul, that's maybe observable that's not necessarily verifiable. These are things that I think you can build metrics around with maybe different confidence or levels of success.

[PAUL HOFFMAN]: Dive into even worse ones.

DUANE WESSELS: And getting even worse, yeah. So application software diversity. This, of course, refers to the software that's actually used serve the root zone: BIND, Knot, NSD, and so on. Routing software.

Lastly – maybe not lastly, but the last one on the screen here – is human diversity, which talks about how we don't want individual persons to be single points of failure. We want service to be operated by teams and to be designed so that, if a key person departs an organization, it doesn't cause too much damage or that one person isn't able to go rogue and do damage to a system.

Is that the last one? No.

UNIDENTIFIED MALE: No.

UNIDENTIFIED MALE: No. [inaudible] documentation.

DUANE WESSELS: Okay.

UNIDENTIFIED MALE: Access—

DUANE WESSELS: Access segregation. That's sort of what I was just talking about, how ideally you want to have segmented access from different staff members so that no one person has too much control over too much of a structure.

Section 3.5 is about documentation and references, again, RSSAC001, how certain procedures and policies should be documented and available. [Purchasing] attack recovery talks about unplanned outages.

When Anycast is utilized, the expectation is that routes are withdrawn from the unavailable sites. Disaster recovery. Again, that's something that's referenced in 001. Backup plans. There's an expectation that operators should have a NOC (Network Operation Center) and should well document how that NOC can be reached and number of staff and so on.

I think I remember this section being a little bit when we were working on this document. It talks about emergency response teams and interaction without outside parties. Establish relationships with – certs are advantageous, obviously, so that operators are well aware of what's going on in the industry and incidents that happen.

All right. Scroll down, Steve.

So the candidate operator should provide sample data, including RSSAC002. The reason for this is because we spent a lot of time on RSSAC002 and getting all the operators to the point where they're providing that. We want to make sure that any future operators also have the ability to do that at the start, going on.

We would expect operators to have their own webpage describing their service and contact details and so on. Then, lastly, this document talks about some kind of evaluation period where a potential operator would be given some amount of time to prove that they're up to the task and maybe possibly even given actual archived data from the root operators as sort of a testing period.

UNIDENTIFIED MALE: [That's it].

DUANE WESSELS: Okay. So we made it to the end of this document. Again, going through this, I think there are some things that are clearly relevant to the Metrics Work Party, some things not. So as we go through the work, we can pick those out and make sure that they have metrics that can be externally verifiable.

Any comments from anyone else about the contents of 024?

RUSS MUNDY: One though that I just had as you were reviewing what was there is there certainly are pieces that make no tie directly into metrics, but in fact tie into potentially other aspects of 037, if 037 continues to go forward, because there's a lot of pieces that have to relate to each other, and some of them from this document particularly really well to 037.

DUANE WESSELS: All right. So the last thing that we wanted to go through – we've got, what, 15 minutes left or so?

RUSS MUNDY: 15. Well, it's [10:15]

UNIDENTIFIED MALE: [inaudible]

DUANE WESSELS: Oh, we have 30 minutes. Okay. Sorry. That's good. So what Steve has put up on the screen here is from the Applicant Guidebook for TLDs. This is registry performance – can you just scroll back? I missed that – specifications. Okay.

So I have to admit, I'm not really familiar with this document. This is a list of definitions, which then feed into this – this is probably the good stuff here, right, Steve? This is the service level agreement matrix.

So let's talk through some of these. DNS service availability. Zero minutes downtime on a monthly basis. Is downtime clearly defined in this document?

BRAD VERD: If you scroll past this – so this is obviously for—

DUANE WESSELS: Executive summary?

BRAD VERD: For a registry which includes both the registry and the resolution aspect of it. Here – actually, I think right below that availability matrix it talked about all the different measurements.

Can you go up, please?

[STEVE SHENG]: Go up?

BRAD VERD: Or down. Whichever. Go back to the availability matrix that showed 100% uptime. So, yeah, just below that if you now scroll.

Here's your DNS metrics that TLD operators have to do for their systems. So this is just another metric that's out there in the community and well-known.

DUANE WESSELS: Right. So, for example, this one says, for the service to be considered available at a particular moment, at least two of the delegated name servers must have successful results from DNS tests. If 51% or more of the DNS testing probes see the service as unavailable, then it will be considered unavailable.

So I think this kind of thing is a good start. It doesn't translate entirely or accurately to the root system, but it's something to think about.

Liman?

LARS-JOHAN LIMAN: I actually worked my way through this in a different role I have as a supplier. If you want to use these metrics, you need to think very carefully because some of them are not well-defined or suitable. For instance, the measurements over TCP are impossible to measure, to gauge.

So I would argue that some of these are not good metrics for DNS in general and therefore we should do better when we do the specification for the root service.

DUANE WESSELS: Would you say that they're not well-defined or that they're not just, as a concept, not good metrics.

PAUL HOFFMAN: They're wrong.

LARS-JOHAN LIMAN: Paul says they're plain wrong. I'd argue that I can see why you want to measure these things, but when you drill down to the actual hard iron, it's actually very difficult to perform these metrics.

So, either way, we need to come up with a measure them and agree that that is the way to measure it. Or we should exclude them or find some other metric that gives us the information that we want.

But be careful when you deal with this Applicant Guidebook because, in my mind, it's not a very good document for this purpose.

BRAD VERD: Again, I wear different hats. In the same [world] you describe, I completely agree with you as far as – trying to implement some of these monitors is challenging, let's say. But I don't think anybody is suggesting that we cut and paste these and create these as our metrics. Nobody is saying that.

But I think, again, if we can scroll back down to the … DNS availability seems like a reasonable metric. DNS name server availability seems like a reasonable metric. How we define that I think is what we need to figure out here. If it's defined wrong, that's fine. We don't need to use that, but it seems like it's something that we should maybe starting with. And the community would understand that.

LARS-JOHAN LIMAN: Fully agreed.

RUSS MUNDY: I agree with what Brad and Liman were just pointing out, that getting an effective measurement on some of these is very difficult. But one of the things that I've wanted to at least raise for the work party to think about is, in terms of defining the metrics – we talked about them being measurable – where and how do we expect these measurements to occur? Is that part of the job of this work party to describe in some manner the measurement mechanisms that we expect to see in place to be used to measure the metrics?

LARS-JOHAN LIMAN: I would say yes. That's my understanding. Or at least suggest methods. It may turn out, as usually, if you design things on paper, they don't really necessarily work that well in reality. But at least a proposed way to measure …

I think we should be prepared to revise the document fairly quickly when we start to deploy measurements when we see whether they work well or not. And we should have a lightweight for doing so. [I'm riding] on experience from the IANA statistics and the CSC because we realized, [with] these two things, that metrics aren't always well-defined from the start. And if you have an extremely heavyweight process to change them, you are in trouble, so to speak.

UNIDENTIFIED MALE: I think it's going to be very valuable to remember what these metrics were intended to be used for. This is like a contract doesn't matter when everybody's happy. This is going to be brought out when somebody is perceived to have failed. By some metric, it will be this metric.

And it will serve us very well if the failed party is able to say, "Wow. I failed. I saw that coming. I can do the math in my head. So can you. We all agree. Now what will we do?" So if you make it one of those tough ones, then there's going to be all kinds of argument over, "No, I didn't," and, "Gravity shifted that day."

So having something that is easy for everyone to recognize and reproduce – and even you can think about a third party brought in to ask if it's fail. This might have to be something that's logically simple for a non-technical third party acting as a judge to come in and grasp.

[KAVEH RANJBAR]: Actually, yes, I was thinking about the same. When some of the metrics – for example, the operating system and things like that – you said are not easy to measure, maybe technically and in live operations they are not. But when things happen and if there's a dispute – because these things will be referred in contracts and all of that. So you never know. There are legal disputes and auditors will come in. Then they look at every single thing.

So, first of all, clarity helps. Actually, in the same sense, I love that we have references, for example, to RFCs, but we also have to think about the control because, when you refer to RFC, for example – depending also on the type of RFC – that means IETF will have control over that because they can obsolete that and come up with new ones, which means there might be measures, others, without at least having to ask you, you can almost participate. But they don't have to ask you and they can actually override them.

So I think we should be conscious about them. I'm not saying it's bad. I'm just saying we should be conscious about that and make sure [inaudible] about references and the clarity.

DUANE WESSELS: Liman, a question for you since you've expressed some experience with these. Is there a contracted party who performs these measurements? Can we get any advice out of them? Or is that a good place to go for advice or is that …

LARS-JOHAN LIMAN: That's probably a good idea. And I have to be close friends with people because the contracted [parting] question is the operators of the Swedish top-level domain, .se. Or a different part of that organization.

So I would be quite willing to at least establish contacts with them and try to get some information back from them. I think that's a very good idea.

DUANE WESSELS: Okay. Yeah?

[PAUL HOFFMAN]: So going back to what you just said about the IETF possibly controlling some things, I just did a quick look. The Benchmarking Working Group had attempted to do DNS benchmarks. I just looked through and, unless this page is wrong, they've failed because there isn't any RFC with that name in it. So I remember this happening. This was like a decade ago.

So not only do we need to be cognizant that the IETF might change some of the requirements, RFC is underneath those. The Benchmark Working Group, after two decades, is still into existence. So that might also come into play.

From my experience with them, especially wearing a previous hat where I ran the VPN Consortium – so we did IPsec testing – that would be horrible for us. That would be a very, very dangerous thing for us to have, them saying, "We want to step in and help here."

I think it would be much better, if we were expecting a third party to do this, that we say upfront we are good at picking a third party, and even if we're not, somebody else in the ICANN realm would be. We don't need this from the IETF.

[KAVEH RANJBAR]: I agree. I think it should be a conscious choice. That's what I'm saying from RSSAC's [inaudible].

DUANE WESSELS: All right. So I guess let's take a little bit more of a look at some of these metrics that are already defined for TLD operations. So the first one was DNS service availability, which is different than DNS name server availability for some reason, although they look sort of similar in their descriptions.

LARS-JOHAN LIMAN: Excuse. Isn't "service" just the combined set of servers? Then it's also per individual server or server cluster or IP address or some definition thereof.

BRAD VERD: The way I read that is the root service and the root name server. Right? That's …

DUANE WESSELS: Okay. Yeah. So there's a metric around resolution round trip time for UDP, for TCP. The RRT is five times greater than – whoops, I wasn't done, Steve. Can you go back? If the RRT is five times greater than the time specified in the SLR, the RRT will be considered undefined. Okay. So that's like a time-out.

Can you go back up to the table? Where … Probably farther … Okay. So in this case, UDP resolution round trip time is 500 milliseconds, and TCP is 1,500 milliseconds. Okay.

All right.

[PAUL HOFFMAN]: Duane, can I? I know this interesting but this is somewhat not relevant to what we are doing here in that. There was a requirement, not just an expectation, that all of the TLD operators would have multiple IP addresses and that each IP address might be Anycast.

Going back to Liman's question about how was this tested, that drove the testers insane. So what I'm assuming we have here is that a root server operator is expected to have one IPv4 address and one IPv6 address. So some of these tests were explicitly like, "Oh, even if one of your IPv4 addresses is down as long as you still meet these." And that wouldn't really work for a root server operator.

DUANE WESSELS: But it works for the system as a whole.

[PAUL HOFFMAN]: Yes, for the system as a whole. Right. Which is inherently 13 addresses – or, I'm sorry. 13 [root servers].

UNIDENTIFIED MALE: Why isn't that just fractal? If the 215 units of a letter work as that whole similarly to the way that all 13 work as a whole, why isn't that a valid test. This sounds like, "Shut off 92 of your machines and see if the 93rd worked." But why wouldn't you be testing it as a whole?

I mean, conceptually, we're so all about, "N minus a pretty big number, we're still functional," so why test N?

UNIDENTIFIED MALE: [inaudible]

DUANE WESSELS: All right. Let's go back to the description list here. Can you go down some more?

So this update time I think is not applicable, probably, for us. DNS tests. So this document does describe in some aspects the number of probes and it looks like probe locations and things like that, so I'm assuming we would want to do similarly, although – yeah. Maybe delay the specifics for future discussions.

Whoops. All right. Can you scroll down?

Ugh. Yeah, I can see where this gets ugly.

FRED BAKER: It actually got ugly much earlier but I'm glad you see it now.

DUANE WESSELS: Distribution of UDP and TCP queries.

UNIDENTIFIED MALE: [inaudible]

DUANE WESSELS: Now what is RDDS? I missed that before?

LARS-JOHAN LIMAN: It's the WHOIS service.

DUANE WESSELS: That's the WHOIS – okay, okay. Right.

UNIDENTIFIED MALE: [So] it doesn't apply.

DUANE WESSELS: So that's not applicable to us. Okay. Then I see there's a section on EPP. Okay.

[PAUL HOFFMAN]: Not applicable to us, right?

DUANE WESSELS: Not applicable.

[PAUL HOFFMAN]: Just checking.

[BRAD VERD]: [inaudible].

DUANE WESSELS: Yeah. Okay. All right, so I think that's going to be good input. We'll do our best I guess to not fall into the same traps that we see are here. But we'll do what we can.

So that covers all the existing metrics and documents that I wanted to go over. Was there any last discussion about that before we wrap up and talk about next steps for the work party?

Yeah?

RUSS MUNDY: I wanted to ask the work party if they knew of any documents that we missed that might exist. Does anybody have any that they know about?

Well, if you do something of something later on, please send it to the mail list. We'll certainly look at it and consider it because we don't want to do anymore than we have to. Yeah?

NAELA SARRAS: Hi. Thanks, Russ. So the only thing I was thinking of when you were doing the technical requirements for new GEs from the Applicant Guidebook is the IANA minimal technical requirements for TLDs. So, separate from that process that the gTLDs have to go through when they come in for delegation or when they do changes on their TLDs, they have to meet specific technical requirements on their name servers.

So that's what I was thinking of, that list, which has the same, pretty much – [reachable] by TCP, UDP, at minimum two name servers, diversity, etc.

So that's the only one I could think of. And I can send you the link if you want, you guys.

UNIDENTIFIED MALE: I think it's an interesting start on a very complex topic, and I think it would really be helpful to repeat one point and make another one, which is to be looking at this from the point of view of the group of us sitting in a circle with one sad member knowing, "Wow. I can't believe how I have continued to fail that well-understood metric we put together," rather than, "There's no way [we've] missed 13 points. The test was wrong. We forgot to allocate something."

I think thinking that through in advance is going to be how – and the second to the degree that we're doing this as system tests rather than component tests I think is important because the heat of the line cards nobody cares about. But the fact that something in the cloud out there did the right thing in the right period of time is important.

It might look messy or unengineeringly, but—

BRAD VERD: Isn't that contrary to what you said earlier? "Why are we testing N?" That's the exact reason we would test N, right?

UNIDENTIFIED MALE: That's what I meant. Rather than peeling off one and testing—

BRAD VERD: Oh, okay. Because that's not what I heard. So I'm sorry.

UNIDENTIFIED MALE: I'm sorry then. What I meant was you have an entity that works as a system and should be tested as such. To arbitrarily pull pieces out and test them I believe is more complexity and misunderstands the system.

I wouldn't risk my life on any arbitrary F-root node, but I would on the whole system still being up.

BRAD VERD: Doesn't every F-root node need to be accountable for the root? If it's –

UNIDENTIFIED MALE: No. It needs to know that it's in a place where, if it fails, somebody else can take care of it.

BRAD VERD: Let me … Okay, well, maybe we disagree. Agree to disagree on that one then.

UNIDENTIFIED MALE: Okay.

UNIDENTIFIED MALE: One thing I wanted to also specify is, in addition to the document, we might have tools or implementation we can also look at and see what the measures are. That might be useful.

But just one clarification. So when you were mentioning we should not measure each node within the system, is the system still the RSO or the 13 RSOs?

JEFF OSBORN: [I'm at] the RSO.

UNIDENTIFIED MALE: Okay.

JEFF OSBORN: I'm at the RSO.

BRAD VERD: Thank you for that clarification. Again, that's not what I had heard.

WES HARDAKER: Wes Hardaker, ISI. So one of the other things that we might want to look at – this is a great list of starting documentation for stuff that's been published. But one of the other things that just occurred to me is there's an awful lot of system and tools out there that measure various aspects of DNS functionality, things like DNSMON, ATLAS, and ThousandEyes and stuff like that.

We may want to go scroll through those. Unfortunately, a lot of those don't really give a pass/fail type, but they at least record metrics. But we want to make sure we're not missing something there.

DUANE WESSELS: Okay. So I'll close out this topic. Let's wrap up with the schedule and stuff. So our plan is to have this work party meet every two weeks, at least until the RSSAC workshop happens in April. Our next meeting would be at the IETF in Prague. I forget exactly which day is it. Monday?

[RUSS MUNDY]: No. Wednesday. 7:45.

DUANE WESSELS: So Wednesday in Prague we'll have a meeting. I encourage everyone to attend.

Also, while I'm thinking about it, a reminder that, if you want to participate in this work party, if you want to be aware of when meetings are happening and things like that, you need to be subscribed to the work party list, which is used for administrative details, meeting coordination, and stuff like that.

The discussions about the work party will take place on the broader Caucus list. But if you want to know about meeting invites and things, you need to be subscribed to the work party specific list, so please get in touch with ICANN staff – Steve or someone – to arrange that.

We'll meet in Prague. We'll have one meeting two weeks after that and then we will also have a work party session, I guess, at the ICANN workshop, right? Or we'll have a discussion.

UNIDENTIFIED MALE: When are those [inaudible] on RSSAC?

DUANE WESSELS: It'd be the RSSAC workshop. Thank you.

UNIDENTIFIED MALE: When is that?

DUANE WESSELS: April twentyish?

UNIDENTIFIED MALE: 23rd.

DUANE WESSELS: Yeah. So where we would like to be with this work party by April 23rd is to have some starting points, some starting metrics, for consideration that we will present to the RSSAC at that time and get their feedback on the direction that we're heading with the work party. We expect the root operators to give their input, their feedback, on what we've got so far and tell us, "This looks good, "or, "This doesn't." A [change of horse]. That kind of thing.

FRED BAKER: How do you want metrics to be proposed? Is that a document saying, "I think we should measure [phoo]. Here's why. Here's how you measure it"? What do you want?

DUANE WESSELS: So I think one of the ways that we're going to get that is by this form that Steve's going to create. So work party members can go through the statement of work and talk about how they are interpreting some of these points and describe metrics.

I think I'm going to volunteer some work for Russ and me. I think, Russ, you and I should probably go through these things that we talked about today and pick out some things that we think are appropriate for the work party to work on and talk about things that are not so we'll have a list of metrics to discuss at our next meeting.

But of course, I think at any time, if a work party member has something they think should be on there, then please speak up on the list or [inaudible] to us.

So I guess we've got a few different ways of collecting the list.

RUSS MUNDY: Yeah. Duane and I haven't had a lot of time to talk about this, but I think we're generally in agreement that we want to draw as much as we can from things that already exist, extract pertinent things, get inputs from others, whether they want to go through looking at those same documents or looking at other information that they have or their own measurements or things that they look at to try in this passthrough to get as wide a collection of what might potentially be useful end metrics for the RSS. And starting out with just simply text and short descriptions of what we're talking about.

DUANE WESSELS: All right. I think that's all we wanted to cover today. Happy to yield the rest of my time to Mr. Secretary over here – oh, sorry.

UNIDENTIFIED MALE: [One minute].

[STEVE SHENG]: Okay. The session has ended. The next session starts at 10:30, a 90-minute session to follow up work on the concept paper. Thanks.

[END OF TRANSCRIPTION]