HYDERABAD – RSSAC Public Session
Sunday, November 06, 2016 – 13:45 to 15:00 IST
ICANN57 | Hyderabad, India

UNIDENTIFIED MALE: Good afternoon, this is RSSAC Public Session.

BRAD VERD: Alright, welcome everyone. This is the RSSAC Public Session, so if you're looking for a different topic or a different group, now is the time to exit. I'm Co-Chair of RSSAC. Next to me is Tripti Sinha, my Co-Chair, and we'll be running through this activities update today. So let's just jump right in.

Slide two, here's what we'll be covering today. We'll give you a quick overview of RSSAC. I know I see a lot of faces that I recognize in the room, so this will be a repeat for some of you and new for others. We will hopefully give enough time for you, but we'll have time at the end for questions should there be any.

We'll cover our publications since the last ICANN. Then we'll jump into reports from the RSSAC workshop that have taken place. We will follow that with updates of the current work going on within the caucus, and then we'll end with community interaction.

Really quick, what is RSSAC? I know people have seen this before, but I'm going to read it because it seems to get misinterpreted and moved around. It says, "The role of the Root Server System Advisory Committee (RSSAC) is to advise the ICANN community and the Board on matters relating to the operation, administration, security, and integrity of the Internet's root server system."

That is a very narrow scope, so just driving that home a bit. Let's move on to the next slide, please.

The organization RSSAC – so there's a formal RSSAC committee, which is representatives from the 12 root server operators. For each of the 12, they have identified an alternate as well. And then of course, we have liaisons to different parts within ICANN that are members of the formal committee as well.

Secondly, the RSSAC Caucus, which everybody in the formal committee is a member of. The Caucus is made up of a body of volunteers, of subject matter experts, and they are appointed by the formal committee RSSAC. Next slide, please.

The Caucus. Currently, the numbers: are we have 77 technical experts from a broad range of different areas. What does it take to become a member? You have to provide an SOI – a Statement of Interest – and the Caucus gets credit for all the work that is outputted from this group. All the advice goes through the Caucus, and they get the credit for it.

Purpose. Again, this is a pool of experts –get all the talent together and work on different statement of works that are brought up, and come to consensus on it.

Again, transparency. Who does the work? It's all very clear in our documents of who has contributed, what their expertise is. And any conflicts of interest are states as well.

If there's anybody here who's not a member of the Caucus, there's a quick shameless plug for the membership. Just e-mail rssac-membership@icann.org, and we can get that going for you.

The next Caucus meeting, if you're not aware, is immediately following this meeting in Hall 1. We will try to leave time for questions. Hopefully, we won't run too long and we can cover any questions that come up, but just note that we do have a bit of a hard stop because most people who are here have to make it to the next meeting also.

Alright, that was the overview of RSSAC. Let's jump right into the meat of why we're here. This is the administration update. First off, the committee has worked with staff and we have come up with a new numbering scheme for all of our documentation moving forward.

Every publication will have a number now, and you'll be able to view all publications by date, by document type, and whatnot on our website. So we've gone back and – we haven't renumbered. We've applied numbers to any document that didn't have a number, and then we obviously kept the old number for any document that was already numbered.

If you can go to the next page, here's the ICANN website. You can see the documents here with numbers next to them and how you can sort by number – which everyone has now – by name or by date of published.

This hopefully will alleviate any confusion that may have happened in the past when there was reference to different documents which, as you can see, have some really long titles and things were getting lost in translation.

I think this is great. We'll be much clearer going forward, and it falls in line with other ACs within ICANN. Next page, please.

Transition. Obviously, this is the first meeting post-transition, so we'll cover different things that happened within RSSAC as a result of the transition.

First and foremost, the root zone administrator – also known as NTIA – liaison has been retired from RSSAC. They are no longer with us. The IANA Functions Operator will now come from PTI, and that person has yet to be identified. PTI is working on that, and we should have that shortly.

And I don't know – what does PTI stand for again?

UNIDENTIFIED MALE: Public Technical Identifiers.

BRAD VERD: Public Technical Identifiers is what that acronym means, if you're not familiar.

And then on top of that, on top of those liaisons, we now have two new outward liaisons that have been created. RSSAC has a standing on the customer standing committee, which Lars Liman over here to my left is going to be taking that role. And then on the Root Zone Evolution Review Committee, I'm going to be filling that role.

And then lastly, which is a very big change, at the end of this meeting, Suzanne Woolf, our current liaison to the Board will be stepping down. She's over here to my left. Kaveh is going to be taking over, also to my left.

Kaveh has been elected the new RSSAC liaison to the Board, but I would like to take this opportunity to thank Suzanne again for your years of service. Thank you very much. Next page, please.

Alright, recent publications. There are a bunch of names here. These people are going to be talking through a couple of the documents. Go to the next page.

Okay, so here, I'm just going to really quickly since somebody is going to talk on each of these different documents here that are numbered. I'm not going to spend a lot of time, but I will point out that as you can see, we have added in the new RSSAC number for all of the documents and any reference to documents going forward will be of the RSSAC number, and not the title as it has been in the past. Next page.

This is the client-side reliability statement, and that's Lars.

LARS-JOHAN LIMAN: Yes, hello. I work for a company called Netnod. We operate i.root, and this is a document that we've published in response to a notion that we've been feeling in the community that, sometimes, we've heard rumors that people think that the root servers are somehow different. They are not.

This is a statement that strives to reiterate that the operators of the root servers are committed to serving the IANA zone in the global root namespace. All members, all root server operators, firmly support the statement by the Internet Architecture Board made in RFC 2826 where they clearly state that the operation of the Internet truly hinges on a single unique namespace. And the root server operators fully support that to 100%.

So, there is no question of the fact that all the root servers, all the machines that provide the service answer with DNS responses that contain complete and unmodified data as we receive it from the IANA through the machinery that provisions the servers.

And to further reaffirm this, we also express our strong support for DNSSEC. DNSSEC is there for the clients to be able to verify that nothing has been modified in transit, and that includes the transit from when the zone is signed through the entire propagation mechanism, including the jump from the root server to the resolver machine.

Through that entire process, it's perfectly verifiable using DNSSEC that the data has not been modified. And every server operated by the root server operators provide exactly the same information. For a given query, you will receive exactly the same data in your response.

There is no difference. There is no server that has better data in any way because it's actually absolutely the same data. And this is provided, also, totally worldwide from all the instances – all the [inaudible] instances where we provide a service. They are all identical, and DNSSEC is the means, is the method, to use to validate that the data is actually correct.

This is a short document. It's only five paragraphs.

BRAD VERD: Thank you, Liman. Moving on, the Statement on Availability of a Single Root. Wes.

WES HARDAKER: Hello. I'm the representative for b.root. RSSAC took upon itself to look upon the history and the past in terms of what has happened with outages on the root, and the conclusion was that if any particular single root server was lost, it would not cause any immediate stability issues for the root server system. And the Internet depends on it, in fact.

Occasionally, root servers do sort of have operational problems, and due to the high redundancy in the root server system, nobody notices. It goes on just as it should.

That, combined with caching is the other thing that… The TTLs in the root zone data are actually quite long, and caching will mean that the resolvers out in the world actually maintain that data for a long period of time, such that the load to the root is generally fairly low in the first place.

And the root server system has experienced, over the years, several really large scale Denial of Service types of attacks, and none of these attacks resulted in any end user visible error conditions that anybody is aware of.

BRAD VERD: Great. I will add to that – maybe take away the thunder from anybody's question that certainly leads to the next question, which would be, "How many roots can be down before there's a service issue?" That is something that RSSAC is looking into and I will be taking to the Caucus in the future. Just putting that out there.

Moving on. The Response to the GNSO PDP, Suzanne?

SUZANNE WOOLF: Yes, that would be me. One of the things you get to do when you've served in a role like mine for as long as I have is – everything old is new again and history starts to count.

The GNSO PDP on subsequent gTLD rounds, I actually sent a set of overarching questions to all the SOs and ACs and other constituencies and stakeholders (so the questions probably looked kind of familiar to most people in the community) basically about differentiation of new gTLDs and how a new round should be structured, that kind of thing.

RSSAC, actually our first input was to state that we have to accommodate whatever the community decides, so we stay away from policy questions. But we did want to give them a couple of points.

First of all, future plans for more top level domains. If they're consistent with the past expansion program, the RSSAC does not foresee any technical issues with further expansion of the root. That’s actually a reference to a bunch of work that was done several years ago as part of the previous new gTLD round to basically do a baseline assessment of how much changes that might look large – how much difference those would make to the functioning of the root and the root servers.

Basically, given the conditions that were discussed and that eventually came out of the new gTLD program, the answer is none. No impact anticipated, and none so far seen.

The other piece of the advice that we sent, though, was to suggest that the root zone management partners and root server operators implement coordination specifically so that root server operators can work with ICANN in the event to identify and deal with any stress, should that occur on the root name service.

And again, that’s just sort of an emphasis on conservatism there because there is a great deal of coordination that goes on between the root server operators and others in the community. But we wanted our input to include that people should consider whether there was any further coordination or any further thought to that that might be needed as part of new gTLDs.

But we didn't feel that there was cause for concern as long as future expansion were consistent with the past. Thanks.

BRAD VERD: Great. Thank you, Suzanne. Next document is The History of the Root Server System which Tripti will speak on.

TRIPTI SINHA: About 13 months ago, the RSSAC decided to begin more focused discussions on questions that needed to be answered, and the questions focused around evolution, accountability, and continuity. So we embarked on what we called workshops.

In support of that work, we decided the only way we could have a well thought out discussion on evolution is to understand how we got here in the first place, so we documented what we call the history of the root server system. All the operators got together and put together their history.

We have compiled this in a document that is labeled exactly that. It has a chronological listing of the history of how we started from a simple file to what we are today, and it's a very good read. We just approved the document, and it will be online this evening, I believe. So please go to our website and you will find the document there. It's a very good piece of history. Thank you.

BRAD VERD: Thank you, Tripti. I will echo her comments that it is a fabulous read. For anybody interested in root servers and how we got to where we are today, this document explains that in very good detail. So please take a look at that. That will be on the website.

Lastly, the Key Technical Elements Document. Lars, are you going to be taking this?

LARS-JOHAN LIMAN: I think that’s where we ended up, yes. The Key Technical Elements Document is actually something that comes out of our second workshop that we held in May where we do look at the kind of evolution of the root server system.

We realized that sometime there will probably be a reason to evaluate potential root server operators, and in order to do that, there are many various things you need to evaluate when you're looking at the alternate root server operators.

The technical elements are fairly easy compared to the others to define. There are political elements. There are financial elements. There are all kinds of various elements. We have to divide this into the various sub-problems, and for each sub-problem, we need to identify who to work with to find the right solution for that problem.

But for the technical elements, we find that RSSAC and the RSSAC Caucus are a very good point for writing a document that defines these technical elements.

It's very important to realize that this document which starts from the two documents that we already have – we have the protocol specification for root server operation, which is in RFC 7720. This is a document that’s produced by the IETF, not by ICANN.

And we have the operational requirements which are defined in RSSAC 001, which is produced within the RSSAC and RSSAC Caucus. Those two are very narrow and they form requirements.

On top of that, there are lots of other technical elements that you need to evaluate. They are not necessarily requirements, but they are things that you need to look at when you evaluate a potential root server operator.

This is an expansion of RSSAC 001 and RFC 7720, and this lists many different types of elements from system design to look at experience and networking properties, look at diversity not only within a certain operator but also in relation to the other operators.

There was some documentation of data and measurements, participating in work with other root server operators – many technical parts that we need to look at and compare when we have multiple potential root server operators. They are not requirements, but they are elements that we need to look at and evaluate. That's it.

BRAD VERD: Thank you, Liman. Again, that document should be on the website this evening, so you can look at it tomorrow. Moving on, we're jumping into our third section about reports on the workshops, and I'll turn that over to Tripti to cover.

TRIPTI SINHA: As I mentioned earlier, we have embarked on workshops. We've had three of them, so this slide speaks to the updates from workshop 2. We held that this past May in Western Virginia, and we took our original themes which were evolution, continuity, and accountability and we began – the metaphor we used is peeling the onion – let's do a deep diver, peel the onion.

For this workshop, we took the approach of looking at the architecture of the system, continued focus on evolution, and what we call reinventing RSSAC.

So, the outcomes of the workshop resulted with three publications which we just went over, so I won't go over that again. It's just the Statement on Client Reliability, The Impact on Unavailability of a Single Root Server – impact on the service, that is – and of course, the technical elements that Liman just went over.

Those were the three publications that were the outcome of that workshop

Reinventing RSSAC. There has been confusion over the years as to how we can approach the root server operators. The RSSAC, as you know, as has been emphasized, we are an advisory committee. That is our mandate. We advise the ICANN community. We don't put on our operator hat. So we've decided we will be the front door to the operators.

So should the community need to reach the operators, you can approach RSSAC and we will then disseminate the information to either all operators or some operators, or an operator depending on what you're trying to communicate. So, those were the outcomes of workshop 2.

Next slide, please. We just concluded workshop 3, literally two weeks ago, and believe it or not, we have a report that we approved which will also go online this evening. So tomorrow morning you should have that available.

For this one, the synopsis is we looked at the entire system from a 50,000 foot abstracted level. The term we're using is mind mapping. We looked at this model of what it looks like today and where we should possibly go, and we put together concepts and functions and so forth, and we're building a mind map.

I will very quickly go over what we've done thus far. This is all available in the report that will be published this evening. The first thing we've done is put together what we're calling a lexicon, because in our own internal deliberations, we were using words slightly differently.

Instance versus node versus server. So we decided the only way we could have any meaningful dialog was to at least agree on what we all meant. That document will also become public when we're comfortable with it and approved it ourselves.

So we put together a list of words that we all rallied around in the context of root DNS services, and then we moved on to a higher level of abstraction of what the system looks like, agreed on that. And then delved a little deeper on topics such as who do we empower, who do we enable? We're building an accountability chain – or an empowerment chain rather – of who's enabled and how do we interact with them.

Then, we went on to the discussion of finance, how we're funding this operation. It's a massive operation, operated by 12 entities. Clearly, finances are important. They need to be important, and we'll keep our eye on that as well.

And from there, moved on to designation and removal. We don't have a process today. It has yet to be determined, but it's something that needs to happen and we will play a role in defining some of those processes.

Then, there needs to be a process for accountability and audit. Who are we accountable to and what are we accountable for? We're looking to define those. And one of the elements under accountability would be to be audited against something. It could be financial audits, technical audits, so on and so forth.

Then we went over technical elements, and Liman just gave you a flavor of the kind of things we're looking at in terms of technical elements of operators.

And we looked at how do we communicate with each other and interact with each other. Today, we've got an informal organization called root ops. Should there be an association? We don't know yet. We have no idea, but we're just keeping our eyes very open. We're having these discussions. We'll put together structures if we feel that should be the case.

Then we're looking at transparency, and Wes will talk a lot about how we've made quite a strong effort to be transparent with the community, and what more do we need to do.

We concluded that we've made some good progress. We're at what we call the inflection point where we do need to evolve for the future in support of posterity – looking at where we are today, this global service that serves billions of people and more billions to come.

So I urge you once again to read the report. It's will be online this evening. It's going through the ICANN process to be posted. Thank you.

UNIDENTIFIED MALE: [inaudible]

BRAD VERD: This was the next slide, but you've already talked through it. This was –

TRIPTI SINHA: Oh, I didn't go on? Sorry.

BRAD VERD: That’s alright. So again, that workshop report should be available tomorrow for everybody to read.

And now, let's jump into current work that is underway. We'll start with the root server naming scheme. I'll give the update on that.

This, in conjunction with the workshops, is what basically spurred on the history document. Basically, there was a discussion that sat down and said, "Let's stop and rethink everything. Is the way that we actually name the root servers today the correct way of doing it?"

We kind of wanted to ask that question, and that’s where this work party started, and work is continuing on this. It's still underway. You can see from the first bullet here, the first thing was to document the history So that’s been done in our history document.

Then the second bullet point there is to consider any and all changes. A number of different options that they’ve looked at as to where the actual names reside and where should they reside going forward. And then obviously, consider any impact to any changes that might happen, specifically around the priming response given now that everything is over DNSSEC.

And then lastly, for them to perform a risk analysis of whatever their recommendation is and provide that to the operators. This work party is nearing completion. They've been underway for a number of months now, and we hope to have this here in the very near future.

Next one is a current work item, and Kaveh is going to speak to this one because he is the work party leader on this. So he's the main guy here.

KAVEH RANJBAR: Hello. Yes, this idea came after, actually, workshop 2 that we as the participants all learned that most of us run Anycast networks – 12 out of 13 actually run Anycast networks. And all of us know our own networks, but we figured out we don't know the full picture.

Like, "Okay, there are 12 operators running this Anycast network. And what are the effects for – if someone wants to add a node, how it affects the big picture? To be able to understand that, we came up with a few questions which we decided to take to the Caucus.

The idea is, for example, should there be a maximum latency for [inaudible] parties defined. And if there should be, what is that number? And again, if there is a number, how you measure it? Do you measure it through the Anycast or through a single root server instance?

Another one is, what happens if you want to add more topologically diverse locations. Going to different locations is more important or going to different networks, for example. All of these questions which we don't have factual-based answers for.

Another one is, should we as root operators – as I mentioned, we run our own Anycast networks, but should we coordinate or shouldn't we? Should we allow this to grow naturally?

And finally, is there any risk to have all root operators as Anycast network or do we still need Unicast operators or not?

There is a work party being formed. The call went out, I think three or four weeks ago, and we have a few volunteers. If you're interested in this work, please –

If you're a member of Caucus, it's easy. If not, join the second session. You will learn how to become a member of Caucus and contribute. And the first calls will be a few weeks after the end of this meeting, so this is quite recent work party which we have started.

BRAD VERD: Thank you, Kaveh. Great progress there. Moving on, Community Interaction. Wes, are you going to grab this one?

WES HARDAKER: Absolutely. One of the things that both RSSAC and the individual root operators have tried to do over the last five years is increase the level of transparency so that the global community has a better view into the inner workings of the root server system as a whole.

We came up with a list of things that we have done, and we actually wrote them down so you can see them. These are sort of the two lists in terms of what we have done to increase transparency.

On the RSSAC side of things, we have established a Caucus, and as Brad mentioned earlier, we have 77 members and it's growing. If you have a technical background in DNS and have input to provide, please join because the more the merrier. And the more diverse opinions we get, the better the Caucus will be – as well as the RSSAC documents that come out of the Caucus.

We publish minutes and workshop reports, of which we've talked about a couple just recently. And one new one will be online later tonight, early tomorrow morning.

We've published the RSSAC and Caucus calendars on when meetings are taking place and what's going on. We have meetings with other ICANN community groups. And if you are a member of an ICANN community group that wants to see more interaction or get a tutorial for what the root server system is – and we feel that there are a number of communities out there that may not be as well informed as you might like them to be – let us know.

We've held a number of tutorials now on what the root server system is, and we're trying to do those twice during every ICANN – and possibly some other external bodies, both actually within ICANN and without.

We have liaison relationships with a growing number of both in-ICANN and out-of-ICANN bodies. And then we have published our operational procedures as RSSAC 000 as the primary core document to RSSAC itself.

On the root server operator side of things, they published minute meetings from any meetings where they have gathered together. They are all publishing RSSAC 002 statistics. This is actually a huge one. You can go dive down into the data of what each root operator is seeing in terms of various types of traffic. RSSAC 002 is the document that defines – if you want to look at the statistics that are published, how to get them and what each statistic means.

All the root server operators have representatives within RSSAC. They have a public webpage now, which is rootservers.org. Most of them, I think – all of them? – have individual webpages as well where you can look up information about their individual operations, and they all have public letters associated with IANA. Or some of them do, no tall of them. Excuse me.

And then they have collaborative reports on major events, so large scale attacks that have been seen across the entire root zone. Documents have been written to describe their view from their perspective.

And then, again, as Tripti or Brad mentioned earlier, RSSAC has offered to be sort of the front door into questions concerning the root server system, so if you have a technical question about how the root server system operates, you can mail the Chairs.

I think there's an e-mail address on the next slide, so why don't we go on? So, where you can send questions about the root server system – that e-mail address is one place to send those questions, and we'll make sure that they get answered.

One thing that we do wonder about transparency is how many people are actually aware of that big, long list. I suspect that many of you – a few of those bold items are quite new, and maybe some of them you knew about.

We're always looking for what's missing from those lists. What else do you want to see us be more transparent about from either group of people? And then how can we further improve that? We are very much looking for feedback in that area.

Please, you can e-mail the e-mail address down there and put “Transparency” or something in the subject line so we have a clue about what it's about. I don't know. Brad, did you want to take anybody from the room at the moment?

BRAD VERD: We can take questions now on this, or we can… Actually, we've just got one more slide. Let's just finish up and then we'll open it up for questions. Last slide.

Obviously, this is where all of our information is. The RSSAC webpage is right there at the top. Our publications, there's information on the Caucus there. And then, obviously, there’s the e-mail on how to join the caucus as well.

With that, that will conclude our RSSAC activities update, and we will move to questions about transparency or any of the other items that we've covered here today. There are a few empty spots here on the table. If you have a question, please step up and start a queue because we don’t have a mic anywhere.

UNIDENTIFIED FEMALE: Actually, I think Steve was [inaudible]

WES HARDAKER: Yes, hopefully we'll have a roving mic. We could use the one at the end if not.

BRAD VERD: Go ahead.

[DAJESH DIN]: I'm [Dajesh Din]. I'm new to RSSAC procedures. Is there any change which has happened or happening post-transition? I just wanted to know that. Thank you.

BRAD VERD: Yes, we covered that back earlier on the slide deck. But really quickly, the root zone administrator liaison which was NTIA has stepped down. That position has been retired. And then there are two new liaisons that have been added, which is the Customer Standing Committee which is part of the transition.

And then also, there's a liaison for RZERC, which is the Root Zone Evolution Review Committee. Those are two new liaison positions that were created as a result of the transition work done within ICANN. That's in the document.

[DAJESH DIN]: So the approval which was coming from NTIA earlier will come from whom? Any root zone changes?

BRAD VERD: Just to be clear, NTIA authorized changes.

[DAJESH DIN]: Oh, I'm sorry.

BRAD VERD: That’s okay. That’s quite alright.

[DAJESH DIN]: I meant authorization.

BRAD VERD: The authorization is now… That is moved to IANA or PTI now, and PTI will be dealing with all of that.

[DAJESH DIN]: Sorry [inaudible]. Will PTI decide on its own? Who has the authority [inaudible]

BRAD VERD: That is out of scope for RSSAC. I apologize, but that’s not ours to speak to.

[DAJESH DIN]: Okay.

BRAD VERD: Steve, did we get a mic? Great. Any other questions?

K.S. RAJU: Mic?

BRAD VERD: Please identify yourself and your affiliation.

K.S. RAJU: Hi. I’m from ISOC and an ICANN Board member. As part of a new transition [model], RSSAC [model the same], [we work][inaudible] creating for the new zones and root servers for IoT and others, as well as based upon the population also.

Actually, we're looking for connecting the [next] 3.4 billion people for the 20 years. If there are more root zones and root servers [models], it will be good. And also, [segregate] for our separate root server root zones for IoT plans, that would be good. [inaudible] for the [oral] process, actually.

BRAD VERD: I had a hard time following all of that, I'm sorry. I think there was a question in there for how to add more servers?

K.S. RAJU: Yes. We'll add more servers, more extra root servers.

BRAD VERD: But it was specific to IoT? Is that what the question was?

K.S. RAJU: Yes. IoT basis and population basis also – global population.

[DAJESH DIN]: Could I just… He means that now that more devices will come onboard – from 4 billion to 20 billion.

KAVEH RANJBAR: Just to answer one part of it, which was multiple root zones, that's definitely not something for RSSAC, and we definitely do not support it. As ICANN, as RSSAC, and as the community, we only go for one root zone, so no fragmentation in Internet.

I know there are some talks about IoT and maybe a different namespace or whatever, but as long as it's related to Internet, we go for one Internet which is one root zone. For different servers, I think that has been covered.

BRAD VERD: Yes. I think, adding servers, we've talked about in previous meetings and we've talked about here. We are looking into how to add or remove root server operators. That process has not been figured out yet.

We're working through that, we're building the foundation up to that, as you can see based on the publication about the key technical elements of a potential root operator. That's the beginning. I think it's key to note that…

I'm sorry, go ahead. Wes, so you want to jump in?

WES HARDAKER: Yes. One important thing to realize for when it comes to the Internet of Things and other stuff is that they don't send queries directly to most DNS infrastructures, and they're all making use of local Internet service providers – resolvers that are actually doing the brunt of the work.

The scale is very much built up from the resolver side. ISPs may need to deploy more on the resolver side, but the number of queries actually getting to the root shouldn't increase that much because of caching and some of the other things that we've talked about earlier.

The root actually scales very nicely to the Internet of things, unless the number of ISPs actually grow.

UNIDENTIFIED MALE: [inaudible]

WES HARDAKER: To reiterate what he said off the microphone. He is referring somewhat to the Denial of Service attacks that were seen by Internet of Things.

Yes, attacks scale across, I think, the Internet at large needs to be thought about not just within the root, but this is an Internet-wide problem not specific to the root system service itself.

BRAD VERD: Lars?

LARS-JOHAN LIMAN: Yes, and I would like to add that within the current system, there is plenty of room to add servers. We are today at 600 servers, and there is no limitation within the [orders of magnitude] here. So it should be very easy to add instances in many places, in many countries, and maintain a very good service.

BRAD VERD: Just to add to that, Anycast has enabled – it is a technology that we use to scale horizontally, which is the number of over 600 instances worldwide, and that number grows daily. Any other questions?

[GOSE]: Just a query. I'm [Gose] from the government of India. RSSAC has the membership from 12 root service operators, and it has a caucus which has got about 77 technical experts. My query is, is the committee or the caucus working on the technical perimeters, how to increase the number of root servers from 13 to a bigger number? Because I believe there are some technical limitations for increasing that number. That is my first question.

Second question is, who decides? Is it the RSSAC who decides where the instances would be there and where the [mirror] servers would be there. And how is it decided? Thank you.

BRAD VERD: The first question on increasing number of roots, I think that question has been asked numerous times, and work has not officially started within the Caucus. We're building up to that question of, is it more servers or is it less servers? Maybe it's not 13. Maybe it's something fewer. And that is a question that has to be answered and will be answered.

And we are – again, as I stated – building the foundation to get there. So as far as number of servers, we're over 600. And there are a number of roots in this room that are willing to talk to you if you want an instance in your country or in your datacenter or on your network.

The second question on who decides, if you go to workshop 2 report, in there, RSSAC states that there needs to be a process on who decides to add or remove a root server operator and that we should be part of it, but there is no process today.

And again, I don't think you can – our belief is that in order to have that discussion, you have to build the foundation –the technical foundation that leads up to that discussion and that's what we're doing in the Caucus with the work that's happening now.

[GOSE]: I'm not talking of only the root servers. I'm talking of the mirror servers and the instances of these existing root servers. Like, say, if any root servers have got that mirror on their instances. I'm talking about those, not the adding or removing of the root servers as such.

LARS-JOHAN LIMAN: Okay. It's very important to understand that the root service is a technical service that's provided to the entire Internet, and it's provided as a single service. That service is provided from a good number of machines. The number today is roughly 600.

Now, the question could be to ask how many servers do we need out there to provide a good service to the entire Internet population. 600 may not be enough. It could well be that we need more.

The second question is to ask how many operators do we need to handle all this currently 600 servers, and how much machinery do we need behind the scene to provision them from the source of the data all the way out to the end nodes.

That is something that is currently handled by the current root server operators, and I am not aware of any limitations in the current system that we have. There may be a limitation in the

600 – how many end nodes do we have, how many instances do we have that provide service to the Internet.

But in the provisioning system behind, that is not limited and it's very easy to expand if we see a need and if the propagation of the data from the data source to the instances starts to lag behind or not work properly. But that is not the difficult part. It's rather easy to provide that with the tools that are existing. S0 to improve that part is easy.

To improve the 600 is slightly more complicated because we need to get in contact with local datacenters and network providers that can help us locate a node somewhere. But once that node is in place, then the chain back to the data source is easy to provide.

[GOSE]: Thank you.

CHRISTIAN ESSELMAN: Hi. I'm with SIDN, the .nl registry. This is not really a question, but more of a remark. We've been running a few projects that sound kind of related to what's being done in your work party.

We have one project in which we analyzed the DDoS attacks on the root of November 30th, 2015, and perhaps that's useful input for this community so I'd be happy to share that.

We also have an ongoing project on the placement of Anycast instances and the impact on latency, so we've been doing some measurements there as well. Also happy to share that work with you all. So if it helps, then...

BRAD VERD: I think the answer is yes and yes.

WES HARDAKER: Are you a Caucus member today? Because you sound like a perfect candidate.

UNIDENTIFIED MALE: I'm not, no.

BRAD VERD: Please think of joining. Going down here in the front and then the back.

VIKRAM TIWATHIA: Thank you very much. I represent the telecom operators in India from COAI. It's nice to hear about the report on the unavailability of root server. My question is that one of the root server plus the other 600 servers, do we have – and this is my first time at ICANN, so I would be ignorant if it is already there.

So, the question I have is, do we have a schedule or a method of testing our resilience by switching off or conducting drills? Say, you switch off a server and see what the impact is of either the root server or the others. is it practically done, or how does it work?

BRAD VERD: I'll speak to it from an operator point of view. This is more of an operations question than it is an RSSAC policy question here – just to make that very clear – but of those 600 that are listed as on the public webpage for root-servers.org, there are any number of them that go up and down in maintenance at any given time.

So they're taken offline all the time for maintenance. Traffic moves as it should via BGP and they are brought back online. So the answer to your question is, it happens every day.

WIKRAM TIWATHIA: And for the root servers? Because this particular report is about one single root server going down, and with all the Internet traffic going to increase plus security threats of DDoS. So, does it apply there as well?

BRAD VERD: There's a bit of a terminology question there, or there may be a piece to cover, which is the document talks about the unavailability of a single root being one of the 13 unique identifiers in the root zone. Those 13 unique identifiers are comprised of 600 instances around the world. So those instances go up and down all the time.

The document was to address the perception out there that if a single operator or a single unique identifier becomes unavailable, that that's a problem. And the document meant to address it that it's not. Brian, do you want to go first, and then Lars?

BRIAN REID: Yes. I'm part of the F-Root team. I'm also perhaps the only root operator who majored in English in college. Anyhow, I'm the person who's leading the group that is putting together the RSSAC vocabulary, the words and what they mean.

We're going to talk about that some more in the Caucus meeting, but I think I heard you say the root servers and the instances. Did

I hear that correctly? Because there's no difference. A root server is an instance, and an instance is a root server.

The thing that there are 13 of is an administrative entity, and not a box. Anything that plugs into the wall and uses electricity is an instance, and every instance is of equal magnitude and they're all root servers.

BRAD VERD: Thank you, Brian. Lars?

LARS-JOHAN LIMAN: Yes. So what the document tries to describe is what would happen if one root server operator – an organization and all the servers, all the instances that that organization operates – fail for some reason.

This could be a financial bankruptcy that forces that operator to close down all its servers, all its instances. It could be a technical problem in this chain that I described, the provisioning chain, that prohibits the data flow from the data source to the instances.

If there's a major technical problem in that chain, that would have an impact on all the instances operated by that operator. And that would mean a subset of all the 600. How large that subset is varies by operator. In my case, it's roughly 55 out of the 600.

So, the document tries to describe what would happen if, for example, our 55 instances disappear from the network. There will still be around 540 other instances that serve it. But specifically, the IP address that leads to our 54 instances would not work anymore. But there would still be 12 other IP addresses that work.

And then the question, would that situation have an impact on the overall service to the public Internet? That's the document.

BRAD VERD: The only piece I'll add, and then I'll come to this next question, is that going back on…

There is a tutorial that's offered here at the ICANN meetings about root server operations, and we had two of them earlier in the week. There aren't any more this week.

You asked about emergency exercises. There is an exercise that happens on the operation side, and that is, again, on the public webpage offered by the root server operators. There's a report on that.

PAUL: I just have a question on your presentation. You had mentioned you're evaluating conventions for the identities of the 13 root servers, maybe evaluating how they should be named. I just wanted you to clarify that.

BRAD VERD: Yes. As I stated, the question was asked, given that these were named years ago, is it still the right thing to do today? The work party has gone in and come up with a number of different options and they are, right now, evaluating the impact of those different options. Lars, do you want to add something?

LARS-JOHAN LIMAN: Yes. It's an investigation as to should we change the names of the service point identities which are, for instance, i.root-servers.net. Should we change the domain name of that host to something else?

In the way past, the names of the servers – the server identities – was changed in order to create more room in the DNS packets. And that was long ago. Should we look at renaming them again? What could be accomplished? Could something get better if we renamed them again?

BRAD VERD: Yes. The questions were asked, is there a better way to do what we're doing today? And that's what the work party has been looking at.

[SHOA ABADI]: Hello. I'm from India. My first question that I tried to find out on IANA's website, is there any contract between root servers operator and IANA? And if there is not, is there any process on that you're trying to develop?

BRAD VERD: To answer the first question, no. There's no contract. A number of letters have an MoU between…

It's self-written MoUs, but if you go to our client-side reliability statement, all 13 root operators have stated that they will serve the IANA root. So that's a public statement that's out there.

As far as your second question, I think –

UNIDENTIFIED FEMALE: What was the question?

BRAD VERD: The second question – if I can rephrase – is the –

[SHOA ABADI]: My second question is, is there any process you're developing for the contract or after IANA transition you're planning to introduce?

TRIPTI SINHA: Yes. As I said, we are doing studies on evolving the system, and one aspect of that system is accountability. We're putting together an accountability chain, and with that will come auditing and so forth. So this new model will speak to that. We're going down that path.

UNIDENTIFIED MALE: [Do you] want to give it to others? Are there any others before we go?

UNIDENTIFIED MALE: There's a whole bunch.

UNIDENTIFIED MALE: Can we come back to you?

UNIDENTIFIED MALE: Yes, [inaudible]

UNIDENTIFIED MALE: Okay, alright. Let me go with that. [inaudible]

DESSALEGN [YEHUALA]: [inaudible]. My question is about the previous deployment model of the root server, and the plan that we are trying to… [inaudible] if there are any issues that you want to address by changing the exiting deployment model into...

UNIDENTIFIED MALE: We can't hear [inaudible]

UNIDENTIFIED MALE: Turn the microphone this way. It's a direct [inaudible] microphone. [inaudible]

DESSALEGN [YEHUALA]: Okay. If there is a plan to modify or tailor the existing deployment model of root servers with respect to the new changes that have taken place because of the IANA transition? Is there a plan to tailor the root servers deployment with a view, for example, to address single fault point of failure, and probably to give some more autonomy for root zone operators? Is there a planned deployment model change?

BRAD VERD: Is there a planned deployment change based upon the result of the transition?

DESSALEGN [YEHUALA]: The transition, yes. Probably that could be triggered by bottleneck issues that have been discovered so far, and probably if there's a room to give more autonomy for root zone operators because of the new change.

KAVEH RANJBAR Can we go to slide 24, please? I don't know who is managing the slides. First of all, the scope of RSSAC – basically, the change… Removing NTIA didn't cause any change in RSSAC's operation because we just publish the file.

The file is generated somewhere else. The process is out of scope for RSSAC. We have a liaison to RZERC, which is the committee who observes parts of how the file is generated, but we are not part of that. So we just distribute the file.

And as mentioned before, we only do one file. So there's no question of autonomy because there is only one root zone file. It doesn't matter who serves it. Nobody can change it. It's cryptographically signed, so any change is immediately visible and the file will be useless.

So there is no question of autonomy and, as mentioned, the NTIA change didn't have any effect on RSSAC operations, except the change of liaisons. That's all.

LARS-JOHAN LIMAN: And I would like to add that if you look the root server system as a whole, the transition of the IANA stewardship only led to administrative changes to the creation of the content of the database. The database is then published by the root zone maintainer and made available to the root server operators who use this chain of provisioning to make sure that the content of the database, the zone file appears at all the instances.

If that chain breaks somewhere – it could be somewhere near PTI in the beginning, it could be inside VeriSign in the middle, it could be inside a certain root server operator's provisioning system also in the middle, before it hits the instance – the only thing that will happen is a delay of updates.

The instance will continue to respond with the data it has. It may be a bit old, but it has data. And there are only very small changes every day to the root zone. so there can be a delay in the chain for an hour (for two hours, for five hours, for 12 hours) without having any noticeable impact, except for possible a single TLD somewhere.

And that can be important for that TLD, but for the root system as a whole, the chain that copies the zone file is not a crucial part. Normally, it updates within a very short timeframe, but the network will not stop working just because the updates don't reach the end nodes. It will give outdated information, but it will not stop working.

BRAD VERD: Right. Next question.

[ASHISH:] Hi. I represent the government of India. My question is of a technical issue. From what I understand, the root server has a collection of NS records for ccTLDs and gTLDs which a resolver can cache on demand.

Right now, what we are following is a pull model on requirement basis. So just a technical thought came to my mind: is it not possible to have a push sort of model so that as and when an update occurs, root server can push it to all the resolver nameservers with [inaudible] sort of address rather than them pulling it?

BRAD VERD: This design change sounds to me like it would be something you would give to your constituency, which would be then taken to the RZERC for review. Wes, would you like to add?

WES HARDAKER: Yes. Specifically, that's not how the DNS protocol works. We can't create that system because the DNS protocol doesn't work that way. If you wanted to consider changing how information propagated from some sort of system like that, you'd actually need to start at something like the IETF which is designed for protocol changes.

[ASHISH:] [Does it require a] design change?

WES HARDAKER: It requires a major design change, exactly. That's not how the system works at all today, so unfortunately, no, that's not really possible.

And it's questionable if there's technical need or not, too. You have to sort of demonstrate why that would be beneficial and what you're actually gaining from that.

Because the reality is that changes to the root zone frequently don't need to be pushed out on the order of seconds. Unlike many other domains beneath it, the root zone changes don't affect people immediately most of the time.

BRAD VERD: Please speak into the mic.

[ASHISH]: My question was with the [assumption] that the updates are not so very frequent of the [main] root server configuration. The ccTLDs and gTLDs, their nameserver, and these are not that frequent.

WES HARDAKER: Correct. They don't change very frequently.

BRAD VERD: I'm sorry. Suzanne was next here as answering the question.

SUZANNE WOOLF: Yes. Just the other thing that your question brings up is, as my colleagues have said, that's not, basically, how the DNS work. There are lots of other things that people who do various forms of business with DNS would be able to do with your question, but it's important to emphasize that the DNS root is so important that we're actually extremely careful about keeping the operations to the most basic features of the protocol.

One of the commitments that the root server operators make is to field software that's complaint with the DNS standard – not add things onto the protocol that don't have the consensus support of the engineering community as good and safe and useful.

So there’s a great deal of conservatism in how the root zone is operated because it's so important.

DAVEY SONG: Okay, thank you. I’m with the Beijing Institute of Internet. Firstly, I would like to compliment the effort of RSSAC to enhance the transparency during the years. I'm the witness of this effort.

And also, I want to speak to that gentleman. Exactly three years ago, I asked the same question that why the number is 13 or why the 13 cannot expand. But when you approach the questions and find answers and more communicate with the group of people who operate this system, you can find that some are related to the history reasons, some are technical limitations or protocol. But there are tiers.

You can call it [inaudible] of the operation or the system, but it needs some consensus to drive the system forward. So, that’s my sharing with you, and that’s the first point.

Secondly, I want to ask a question. We observed… Some events happened the past year and this year, the DoS attack. So, I also joined the Caucus party about the Anycast technology, how to evaluate and how to develop.

My question is, is the Unicast model the single choice? Are there any other possibilities? Because current Anycast strategy is to isolate the attack traffic so that it will cause the scenario that some area happen to be the area that cannot access to the service because maybe the attack traffic from that area.

It's not easy for the people living there to avoid the attack, so if there are any possibilities to maybe some more solutions or considerations on that. That's my question.

LARS-JOHAN LIMAN: I just have a clarifying question. Are you speaking in the scope of existing standards and protocols, or are you thinking about possible new protocols?

DAVEY SONG: I would like to explore new protocols.

LARS-JOHAN LIMAN: Because there are other possibilities, and I'm sure you have heard about all the possibilities. But the thing is, at the level of – and I'm speaking as a root operator. As Suzanne mentioned, at the level of root, when it comes to real root operators, we only work with very well established protocols.

So I'm fully for exploring the protocols, but that's mostly IETF and forums related. But yes, if there are well established protocols and tested protocols, why not? We are open to that. It's not that we just want to do one thing one way, but there have to be very well established protocols which we know will support the goal of distributing the single root file via consensus.

DAVEY SONG: Thank you.

BRAD VERD: And we've stated that in documents, that we are open to new technology and evolving.

UNIDENTIFIED MALE: Brad. [inaudible]

UNIDENTIFIED MALE: My question is very simple. Don't increasing a monopoly of 13 root servers operated by implementing DNSSEC in the absence of any direct contract with IANA? Because [if you] started implementing DNSSEC–

BRAD VERD: I'm sorry. Can you slow that down a bit and state it?

UNIDENTIFIED MALE: Okay. Don’t we increasing a monopoly of 13 root servers operators by implementing DNSSEC in the absence of any direct contract with IANA? Because I think if we are implementing DNSSEC, it's taking away a technical or…

If anyone wants to create their root server and want to talk to you guys and want to add their IP address, in future, there will be no possibility of – without taking your permission to come into the network at the top level, I guess.

SUZANNE WOOLF: I think it's important – the distinction between the editorial function and the distribution function.

BRAD VERD: Go for it.

SUZANNE WOOLF: Okay. If I understood your question, it's important to distinguish that the root server operators, the people that RSSAC talks with and works with, do not and have not ever been the ones to decide what top level domains go into the root zone. That’s, by definition, the function of IANA. And the root server operators deliberately stay out of those decisions.

DNSSEC actually helps that because what DNSSEC means is that when the root zone management partners sign the zone, even if a root server operator wanted to change the answer to a query – if any one root server operator wanted to change the answer that you got so it was different than what he got – it would be immediately detectable because the checksums wouldn't work. That's what DNSSEC does.

So, if I understood your question correctly, DNSSEC actually helps.

UNIDENTIFIED MALE: One more is [inaudible]

UNIDENTIFIED MALE: I'm sorry. You're got to talk to a mic.

UNIDENTIFIED MALE: Actually, it's not like I don't believe in the root server operators, but one thought came across in my mind when I was questioning myself. After implementing DNSSEC – DNSSEC is a very good thing. It provides security to the end users, like authoritative root server. It's a chain of the keys. It helps you to reach to the right TLD operators and mail servers, everything.

But the thing is that if you implement DNSSEC in that system, it will answer to the only 13 addresses of the root server operators. I think that if I'm not having a correct idea about it, then you can clear it.

WES HARDAKER: Tell me if I'm getting your question correctly. Because DNSSEC actually signs the list of root servers that signs the name server records within the root zone, that that list cannot be modified except by IANA itself.

That’s always sort of been the case. The root hints file distributed by IANA as well as the root zone file has always been solely in control of IANA. That has never been in control of anybody else. Now, if anybody did try to modify the list of 13 name servers, yes, DNSSEC actually allows you to detect any of those changes. But IANA has always been solely in control of that list anyway.

LARS-JOHAN LIMAN: We have time for one more question, the gentleman over there.

DMITRY BURKOV: Good afternoon. I want to remind my colleagues, because about two or three years ago, it was also about [inaudible] of Geoff

Huston [inaudible] on the request. And secondly, a few proposals from Vixie, from Kumari and so on.

Because what I heard here, what you try to develop, is traditional approach. It looks to me that you ignore that you are m6ore and more just a distribution system, or maybe transport system.

Because if those things that [inaudible] was identified by Geoff, it was even two years ago we have approximately 50% even regarding his methodology. Where missed you guy? It was [through] Google and so on. The reality might be worse because a lot of operators, in fact, have a cache of copies and distribute their own. What is the current situation?

And what's also interesting, what do you think about future? What is the role of root server operators? Is it really to answer the results, or just to provide transport?

Because, of course, remember, it looks like three years ago this seminar in Hong Kong when the different [inaudible] draft were discussed and all sold. But I never heard later except [inaudible]. What's happened with all these drafts? What do you think about this? What do you think about strategical vision for the whole system?

WES HARDAKER: Good question. There have been other people who are trying to do work in that area. That is sort of beyond the scope of RSSAC, because RSSAC is about keeping the stability and the resilience of the root servers as they are today.

If other technologies come along and replace the existing system, of which –

Hold on one sec because Warren Kumari, one of the people you were talking about – actually, his draft in question did make it to an RFC. It's 7761 or something like that.

DIMITRY BURKOV: But [inaudible] if you will limit just operator stability or root system stability, it's very significant.

BRAD VERD: Wes, I think what we're after is service continuity, availability of the root service. Okay? Just to be very clear.

DIMITRY BURKOV: Why I stress, stress anyway because on whole, you can't ignore the changing world around us – that you are less and less, but you are a significant part to distribute from master zone. But the whole system can be destabilized different way, no matter how you will be successful.

BRAD VERD: Kaveh?

KAVEH RANJBAR: You're right, but there are new ideas and new standards. But actually, so we will look at – one example was the Anycast document that was shown. And actually, I'm reading from a statement of the work that says, “It's worth noting that solutions that involve locally mirroring a root (RFC 7706, for example) need to be considered or referenced by the work party.”

So we’ll look into those work which has been done and we will try to see if we can incorporate or not. But I will again repeat what Suzanne said, which is, in general, root operations – we are very conservative. But we fully understand, and it's, I think, in even all the [inaudible] of statements that we understand the changing environment, and we have nothing against accepting change.

The only red line we have – and we have stated that – is one single root file. As long as that’s there and we can distribute that safely, there is no other thing we are not open to change.

BRAD VERD: With that, we've run out of time. Again, we have a hard stop and we're moving to the next meeting. In Hall 1, the RSSAC Caucus Meeting. I hope to see you all there. Thank you.

[END OF TRANSCRIPTION]