SAN JUAN – Joint Meeting: OCTO and RSSAC
Tuesday, March 13, 2018 – 13:30 to 15:00 AST
ICANN61 | San Juan, Puerto Rico

TRIPTI SINHA: Do we know who's going to run that?

UNIDENTIFIED MALE: You can.

TRIPTI SINHA: I'll just kick it off. Okay. All right. So, I want to kick off the Joint OCTO RSSAC meeting. Welcome, everyone. We've got a full agenda and I'd like to start by saying there's been a lot of back and forth between OCTO and RSSAC and we've been emotionally riled up to the point where David is so choked with emotion, he's lost his voice. Okay?

So, my apologies if I've given you grief. All is well.

UNIDENTIFIED MALE: I'm used to it by now.

TRIPTI SINHA: All right. So, I think first on the agenda is KSK rollover, right? That's true. Okay. You're absolutely right. This is a public

meeting. Okay. I'll start and then we'll just go clockwise and go around the room and to the back. Tripti Sinha, University of Maryland. Co-Chair of RSSAC.

MAURICIO VERGARA: Mauricio Vergara, ICANN.

DARREN KARA: Darren Kara, ICANN.

BRAD VERD: Brad Verd, Verisign Co-Chair.

LARS-JOHAN LIMAN: Lars-Johan Liman, Netnod, member of RSSAC.

VICKY RISK: Vicky Risk, ISC.

FRED BAKER: Fred Baker, ISC.

JEFF OSBORN: Jeff Osborn, ISC.

NAELA SARRAS: Naela Sarras, IANA Functions Operator and ICANN Staff.

RYAN STEPHENSON: Ryan Stephenson, DoD.

RUSS MUNDY: Russ Mundy, SSAC Liaison to the RSSAC.

JOHN CRAIN: John Crain, ICANN Alternate and RSSAC member of the ICANN OCTO Team and many other hats.

SUZANNE WOOLF: Suzanne Woolf, RSSAC Alternate.

WES HARDAKER: Wes Hardaker, University of Southern California.

UNIDENTIFIED MALE: [inaudible] SSAC Liaison.

CATHY PETERSEN: Cathy Petersen, ICANN Org, Office of the CTO.

MATT LARSON: Matt Larson, ICANN Org, Office of the CTO.

MARIO ALEMAN: Mario Aleman, ICANN Staff Supporting RSSAC.

STEVE SHENG: Steve Sheng, ICANN Staff Supporting RSSAC.

KAVEH RANJBAR: Kaveh Ranjbar, RIPE NCC.

DAVID CONRAD: David Conrad, ICANN CTO and rhinovirus carrier.

TRIPTI SINHA: Could you just quickly say where you're from and announce yourself.

UNIDENTIFIED MALE: [inaudible] from JPRS.

PETER KOCH: Peter Koch, DENIC.

MICHAEL CASADEVALL: Michael Casadevall, Freelancer.

MIKHAIL ANISIMOV: Mikhail Anisimov, .ru registry.

CARLOS ALVAREZ: Carlos Alvarez, ICANN SSR Team.

TRIPTI SINHA: And we've got many others online, as well. Welcome to all of you. And with that said, where's RSSAC? Colleague, go ahead. Over to you.

UNIDENTIFIED MALE: So, I believe we have some slides loading. I think actually to save my voice, I might throw this over to Matt. This is a talk on the KSK rollover.

MATT LARSON: Hello, everyone. I think rather than go through this entire presentation, Steve, could I ask you to advance to Slide 6, please? It's a table. Yeah. I'll just sort of recap. I've given this presentation once already and I think two or three times more, and people have either seen it or will see it, so I won't do the whole thing in the interest of time.

But let me just highlight the schedule. So, very briefly, earlier this year, I guess it was actually late last year that we announced our intention to solicit input from the criteria on input from the community on criteria to proceed with the root KSK roll. And so we asked for input on a mailing list called ksk-rollover@icann.org, which everyone should please subscribe to. Ksk-rollover@icann.org. And we got some comments there that amounted to endorsement to continue the KSK roll but to keep reaching out and doing the media push, and you're getting the CliffsNotes version here in the interest of time.

What we did after that, then, was on February 1st and this is where we pick up on the slide on February 1st, we published a draft plan that called for the KSK roll to occur on October 11, 2018, exactly a year delay, and this then is our hoped for plan for how to proceed from now until October 11th.

An important point, though, is that the plan we publish is a draft plan and we're really serious about that. It is a draft plan. We want community feedback on that, so to that end, there's a public comment period open right now. It's been open since then. It closes on April 2nd. There's a distressing lack of comments so far, so it would be great if we could get more comments. And based on the feedback we get, we will revise the plan as necessary, and the revisions depend on what we hear.

If the revision don't cause us to alter the timeline up here, then this is what we envision mid-April. Then we publish the staff report and public comment, which happens after every public comment and the revised plan. Then there's a board workshop in May and we'd ask that the board ask SSAC to review and then based on discussion with Kaveh – I'm sorry, the slide has not been updated – we would also kindly request RSSAC if they would be willing to review the plan, as well, at that point and provide feedback.

Then in ICANN62, we'll have another session for feedback and then by August 1st, we would hope to receive the SSAC and RSSAC feedback on the plan, make any revisions, publish the final plan in mid-August, and in September, get a board resolution with authorization to proceed on October 11th.

So, that is what we're looking at in terms of scheduling and so I guess the main thing I'd like to do here today in addition to just updating you and asking everyone to please consider making a public comment, is to just I guess semi formally ask if RSSAC would be willing to review the plan and provide feedback, if that's something that the committee could do, we'd be very grateful.

BRAD VERD: I think the answer to your question is yes. We'd love to review the plan and we look forward to it. Is there… I don't have any of those questions about what the impact to the root system might be, if there is any for this, given the numbers you're seeing and data that you're seeing. I don't know, I'm just –

MATT LARSON: I'm sorry. What's the question?

BRAD VERD: Has there been any research or thought on what the potential impact to the root servers, if any? Maybe the answer is none. I'm just curious if the question has been asked. That's all. I know it's been asked in our group, so I thought I would share.

MATT LARSON: Yeah. I think the running assumption is that there won't be any impact to the root server system itself. It'll just keep resolving queries as usual and if it returns a signature over the keyset that somebody doesn't have a trust anchor for, then that person will have a bad day. Some resolvers do get more aggressive when they encounter bad signatures, so it is possible.

BRAD VERD: You mean like a re-query or something like that.

MATT LARSON: That sort of thing, yeah.

BRAD VERD: All right. Any other questions or comment?

UNIDENTIFIED MALE: So, one of the suggestions that came up because there is a logic to this plan and the board resolution will be in May, but we know for a fact that there will be a resolution. We don't know the content of it but we at least know the relevant questions from RSSAC. So, one of the suggestions both for RSSAC and SSAC and it was also discussed with the SSAC I know is to we can be proactive before the resolution. There will be a resolution and then based in that resolution, RSSAC and SSAC will be formally asked to look into that and provide advice, but we can be proactive.

Nothing is stopping us to stop thinking about that from today and even if we have advice, board is more than willing to receive that before issuing that resolution in May. The reason for the May is basically that second drop April should pass, there should be the staff report published. There is [inaudible] period for the board to be able to make a resolution, so there is some

[inaudible] involved, but it shouldn't stop us from actually thinking about the issue, if you choose to.

BRAD VERD: Russ?

RUSS MUNDY: Thanks, Brad. In an effort to try to get a little head start, I know that the information is out there but the plan and the open public comment, it probably would be helpful to RSSAC and I know it would be me for the SSAC liaison to get the current set of pointers to the actual rollover place need the public comment and have either John or Kerry put it on the list so it's there and we can do our public comment responses individuals and also start reading the plan so we have a better idea of what it's going to say when some actual tasking shows up.

MATT LARSON: Okay. I can send an e-mail with that to Brad and Tripti. Okay.

BRAD VERD: Yes. We will disseminate it. Thank you. Any other questions?

TRIPTI SINHA: So, I'm going to move the agenda around a little bit because David needs to go to a different meeting. David recently did a presentation to the board to talk about concerns regarding DDoS attacks to the root server system and we're also talking about hyperlocal roots and so forth, so he's going to do a quick review of that discussion and then I think he needs to bolt and we'll continue on with the other agenda items. Thanks. Over to you, David.

DAVID CONRAD: Thank you. This presentation is a result of a request from the Board Technical Committee and I guess Cherine to talk about what ICANN, the organization, can do with regards to what the board is increasingly seen as an existential threat to the root server system as a whole and to ICANN's role as part of the root server system.

This talk, this presentation and the board paper talks only about what ICANN the organization can do because that's the only thing we really control, and that was reiterated to the board on a number of occasions that this is not talking about the system regardless of what the title says, but says it's specifically about what we can do within the context of ICANN.org to try to address some of the concerns related to denial of service against the roots. Next slide, please.

So, a lot of this was setting the stage to make sure the board members understood that the DDoS threat against pretty much any target on the Internet today is actually real, and made reference to the 1.7-terabit attack against GitHub as a result of the [inaudible] vulnerability, and did point out that it only took 10 minutes to mitigate. However, the real issue there is that if, say, the root service was down for 10 minutes, it wouldn't really matter if that had that much of an impact from the context of the board because it would call into question ICANN's ability to help coordinate the Internet system of unique identifiers and probably result in congressional testimony and generally a bad day for the board members. Nothing else really to talk about on this slide. Next slide.

Provided some information as to what our OCTO's view of what the main causes for this are. It's been pointed out that it isn't really only IoT devices because [inaudible] isn't an IoT-related vulnerability but our concern is driven largely by the proliferation of devices with essentially crappy security and as a result, creating an environment in which it's relatively straightforward to create a denial-of-service vector of essentially unlimited capacity against any sort of infrastructure that does have limited capacity. Part of the contributors to that is the fact that it's still trivial to spoof addresses and amplification is increasingly a concern.

Sort of the takeaway on this was that there isn't an easy way to address this on the supply side and because we can't address it on the supply side, we have to look towards solutions on the receiving slide. Next slide, please.

So, the board was specifically interested in options that ICANN can apply unilaterally, so the ICANN org can apply unilaterally. And those were to expand the L root, both in terms of the L single deployments within hosting organizations as well as L clusters. One additional version, which is talked about in the paper but not so much talked about in the slides is the idea of working with cloud providers and talk a bit more about that later.

Encouraging hyperlocal root server deployments, so basically as I'm sure you're all aware, mirroring the root contents into more closely to the edges either into resolvers or the support authoritative or an authoritative next to a resolver and addressing that via 127.0.0.1 or something like that or even you could even view the L singles as another form of hyperlocal one that's more traditional in deployment.

Research at DNS protocol enhancement some like the hammer draft that Warren put out or the TTL stretching, looking at trying to encourage deployment of things like [inaudible] use and those sorts of things, all aimed at providing improvements in the

protocols to help address denial of service and the risks associated with denial of service, understanding that none of them are silver bullets, that each one will address a small bit, but every little bit helps when you're trying to beat back massive denial of service. And there were a couple of other options that within OCTO we didn't feel were worth pursuing but have been mentioned in numerous places, and we though the board needed to be made aware of that. Next slide.

In terms of expanding L, so it's basically adding more anycast instances, whether it's a small instance in the case of L single or adding/creating new clusters, we currently have three clusters. We're looking at adding additional clusters, that's an expenditure by ICANN, not [trivial] expenditure.

The L singles deployment are the cost, most of the costs are handled by the hosting organizations. There is some administrative overhead that has relatively small amount of expenditure, but this is something that we can do completely unliterally, obviously, working with partners. The point here is that this is not a sustainable solution. We do not believe that we can continue throwing hardware and bandwidth at this problem and stay ahead because of the proliferation of IoT devices and vulnerable systems that are connected at high bandwidth. We just don't think that anyone – not just ICANN – but anyone can win that race. Next slide.

The hyperlocal – I'm sure it's been discussed a number of times here. A form of hyperlocal is 7706, that's not the only one. There are a bunch of others. We within OCTO feel that it is probably the most appropriate midterm solution for the denial-of-service threat as it currently exists. It's not an exclusive. All of these things have to work and play well together moving forward but having the hyperlocal approach means that any denial-of-service attack would be targeted at the entity who has much greater ability to control the input of that attack, presumably going out and turning off their customers that are starting to hammer them.

We feel that this is a scalable and decentralized solution but decentralization does have downsides. It means we lose visibility into what's happening within the context of root service. There have been arguments that it is operationally more fragile. We can have differences of opinion about that. I personally think it's a self-correcting problem if somebody misconfigures something of this nature, they'll have quite a significant incentive to fix it. But we still believe that this is the best medium-term solution for addressing an existential risk to the root servers. Next slide.

There are a bunch of DNS protocol enhancements that are ongoing right now. I won't bore you all with going through them, but they include improving the resiliency of the protocol in the

case of denial-of-service but also changing the underlying transfer protocol to try to get away from the sort of the sucking chest wound that is amplification source spoofed UDP packets.

From ICANN's perspective, the actions that we can take there, just continuing to support the research into these enhancements, funding prototype and pilot implementations where appropriate. The issue there is that to have any impact on the denial-of-service risk, it's sort of a long time horizon, getting the code out and deployed to a point where it'll actually have a meaningful impact on denial of service is undoubtedly quite a long time. Next slide.

The other options I mentioned – so if you sit down and do the math, 13, we learned quite some time back is not a fixed hard number. There are now 26 IP addresses in the priming response, more or less, depending on which server, you get more v4 or v6. We didn't think this was worthwhile pursuing as an active measure against the denial-of-service because it is nonsustainable, just like adding more instances. In fact, it really is just adding more instances. They just happen to have different IP addresses. It doesn't really solve any problem and creates just an incredible incentive to buy popcorn to watch the political fallout.

Similarly, there has been discussion about the idea of whitelisting and filtering at the root servers, so the idea there is you as a large resolver operator, can go to the root servers somehow and say, "Here is my IP addresses. Please prioritize queries from this IP address." And folks who are not prioritized would be the first to be dropped or could be dropped at all times.

This obviously does not address any denial-of-service that's based on volumetric, the ones that are attacking the routers in front of the DNS servers. This wouldn't help at all. It is also likely to be quite politically charged because some would argue that it would go against the idea of the permissionless Internet. In this case, if you're a resolver operator, you have to ask for permission to have your queries answered.

Counterargument to that is well, then those folks can just do hyperlocal, but we just felt that this is a Band-Aid on a much more serious problem and this particular solution, which undoubtedly people will deploy just to protect their own infrastructures from attack, isn't one that is appropriate in ICANN's context. Next slide.

This slide just summarizes the four… There are six – one, two, three, four, five, six, yes, I can count – options that we explored

with the board and I'll end it up to questions at this stage. Or comments. Screams of outrage. Any of the above.

BRAD VERD: Thank you, David. I thought that was very helpful. I think this echoes a lot of what's going on within the root operators and their practice, so this isn't necessarily new. I do have a couple of questions that I'll ask based upon this. So RSSAC – we got a hold of this early, so there were a couple of questions that were brought up and I'm going to channel the group by asking them. So, don't kill the messenger.

First question around hyperlocal, there was a kind of a clarification – it was either needed or asked for in that this doesn't enable localized damage in the event of a DDoS or at least we didn't see that or the discussion was since it's at the resolver, the attacks will be coming directly to the roots, it would certainly keep the resolver working, but it doesn't mitigate the attack to the root server system in –

DAVID CONRAD: Yep, very true.

BRAD VERD: It's really hard to talk about because attacks are so complicated and the attack vectors change all the time, so this addresses one of the attack vectors and it's not… yeah. And I guess some of the discussion was that that needed to be a little bit more clear.

DAVID CONRAD: Sure. Yeah. The focus here was – and that's why I apologize for the title – it's really on resiliency for root service as opposed to the root server system or root servers or anything like that. It's the ability to get the referrals from the root so that people can then query down the tree to get the names resolved. So, the focus here, the interest, really, is on trying to figure out ways of protecting root service. And in that context, hyperlocal would provide root service because the root is brought in, and if you're experiencing a denial-of-service against your resolver, then you, as an ISP, would be able to identify the clients that are beating the crap of it or do RL or the whole variety of things that you can control as the network operator.

The question about denial-of-service against the root in a hyperlocal environment, that gets back to the same problem that we always face is the only real answer we have at this point in time is throwing more capacity and more anycast instances at the problem. That's why I said that none of these are exclusive. We're going to have to do all of them and in the end until we

address the larger problem of ensuring that spoofing [isn't] continues to be a problem or the amplification doesn't continue to be a problem. We're not going to be able to solve all of these problems.

WES HARDAKER: Can I ask a real quick terminology question, David, which is the usage of the word hyperlocal sort of popped up out of the middle of nowhere, and I know you guys have an internal project to working with NL NetLabs and I think ISC and other stuff to get stuff deployed into a codebase. Do you consider hyperlocal to be the generic term for caching locally or is that a codeword for your specific project? And if not, do you have a word for your specific project that can refer to [with this]?

DAVID CONRAD: The term hyperlocal is actually something I believe Steve Crocker had coined to reference – well, actually, I should say this is my interpretation of Steve's term, which is replicating the root zone locally and then serving that root zone locally. There are a variety of ways that can be done. 7706 is one of the ways they're, as I suggested, L singles could also be seen as a hyperlocal approach because you're bringing it, you're bringing the root zone closer to the edge, so it's intended to be a generic term,

meaning generic term that references the concept of replicating [the zone] locally.

With respect to the effort to improve the codebase, I will let you in on a little secret. Internally within ICANN, we don't necessarily communicate all that well together, and I actually do not know what [Terry's] group who's funding the work with ISC and I guess NL NetLabs is actually doing in terms of the hyperlocal, so that's a question you should direct towards [Terry] at this stage.

BRAD VERD: All right. Thank you for the clarification. Is there any other questions? Yes.

UNIDENTIFIED MALE: When fusing hyperlocal, is there a concern that that would in some ways make maintenance and upgrades to the root more fragile? Because then you're dependent on the hyperlocals to download new top-level domains or key rollovers, or if there's large protocol changes, because if some of the larger DNS resolvers say like Google or OpenDNS switch over to hyperlocal and for whatever reason, they get stuck on an older version. Wouldn't that in some ways hamstring ICANN's ability to update the roots?

DAVID CONRAD: So, my view on that is that we have to assume that people are responsible adults. I believe… I don't know if [Warren's] there. I haven't seen him. My understanding is that Google does do hyperlocal implementation that Google will never query out to the root servers. We have to assume that they will follow appropriate timers associated with the SOAs for the root to pull down updated zones and that they won't let the zones expire, neither the signatures or the actual data itself.

Similarly, I think my understanding of the work that [Terry's] team is doing or funding, rather, is to try to make it turnkey so that the people who choose to deploy hyperlocal will not shoot themselves in the foot. This is in some ways similar to sort of what we're experiencing with the KSK rollover. We're finding upwards now of 28 plus percent of resolvers that are reporting 8145 data are showing RFC 2010 only, which suggests that they're going to be having a bad day. However, this is a self-inflicted wound and when they do have a bad day, they will remedy it if they feel the need to. It might be fun to not be able to resolve anything with DNSSEC, so similarly with hyperlocal, if there is a misconfiguration that causes a zone to expire. And I'm not too worried about zone data getting stale because if it does get stale, then the signatures will expire and the server will SERVFAIL.

But if someone is not actually keeping the data up to date, then within a week or so, whatever the timers are, I haven't looked in a while, they will start getting phone calls on their support lines saying, "Why can't I resolve anything?" and it's a localized failure within the service area of that resolver.

UNIDENTIFIED MALE: Okay. Thank you.

WES HARDAKER: I can channel Warren for you really quickly because I had a conversation with him. Google does not regularly keep hyperlocal support actually in their servers. I think they might have the ability to turn on. I can't channel that side of him officially but he did tell me that no, they don't normally do that.

DAVID CONRAD: Hmm. Interesting. Okay.

BRAD VERD: Yeah. I was going to add that we see lots of queries from Google to our roots, so. Jeff was next in the queue here.

JEFF OSBORN: Dave, when you describe the concern of the board as having a feeling this is an existential threat, those of us – you and I have been doing this a long time and there are always risks and always threats and always things you can imagine that are pretty awful – are you concerned this happens to be happening or do you genuinely see a tipping point coming that is in the offing that requires radical action?

DAVID CONRAD: I can give you my opinion but I don't think that's actually relevant. My impression is that there are board members who believe that the proliferation of IoT devices has actually fundamentally changed the game, and that this is an actual threat to the Internet as a whole, and the only area in which ICANN has impact is in the area of root service.

So, in the view of these board members, their view is that we have to at least do what we can do to try to protect the part of the Internet that we have some control over.

BRAD VERD: Russ was first. Russ then Wes. I'm sorry.

RUSS MUNDY: Thanks, David, for a lot of good information here. I think part of the problem that we face is, indeed, something we've talked about both in RSSAC and SSAC for a while, and that is the difficulty in trying to understand how the root server system as a whole is performing.

As part of this effort to examine how to improve things, has there been any additional thought given or additional impetus put into how we might go about doing the measurement of the system as a whole?

DAVID CONRAD: There's been some informal discussions about doing additional studies, additional mechanisms by which we can establish behaviors and gain more information about how the dynamic system is actually operating, but that hasn't… I think… I'm trying to channel board members here. I think their view is that that is just day-to-day business of my team and when they have specific questions that they will then submit those to RSSAC or SSAC. I don't think beyond saying yeah, we need more data, that they've given it a whole lot of thought.

BRAD VERD: Wes.

DAVID CONRAD: Actually, before… I would actually defer to an actual board member here, if he wants to correct my statements at all.

UNIDENTIFIED MALE: I think actually it's a very good summary in board, especially the issue of IoT, which is now the talk of the day, has been brought up multiple times and I think that was a very fair summary of how some of the board members feel. And you also have to understand it's hard. I really appreciate this type of work because it shows what can be done to mitigate those fears, let's say, but how realistic those fears are, that's something, which actually is very hard to argue within ICANN framework, because most of them are IoT and it will be really hard to spend board time to educate them on IoT, which is completely is out of ICANN remit, so there is also that thing in place.

DAVID CONRAD: Yeah, and I will say that in the spectrum of board members, I can state authoritatively that there are people who feel that they're surprised, they don't wake up every morning and the Internet, the root service is completely burned out of existence because of a denial-of-service to the other extreme where it's like Internet is going to die every day. This is [MPEG at 11]. It's not a big deal.

So, part of this exercise, the board was basically looking for proposals that the org could implement in order to address the concerns that some board members see, so that then they can think about the cost/benefit analysis.

BRAD VERD: Michael, you had a question.

MICHAEL CASADEVALL: [inaudible].

BRAD VERD: Wes asked the question. I'm sorry. Should we go back to Wes? Do you have a follow-up?

UNIDENTIFIED MALE: No, I interrupted Wes [inaudible].

BRAD VERD: Oh, I'm sorry.

MICHAEL CASADEVALL: One other thing with hyperlocal, though, is doesn't it just defer the problem? If someone could actually get enough bandwidth from IoT devices to knock off the entire root server, it is

conceivable that they could keep the servers down long enough to time out the keys and then we're right back to where we started, unless you plan to implement filtering for the hyperlocal updates, which would at least allow the hyperlocal roots to keep updating while the rest of the root servers are on fire.

UNIDENTIFIED MALE: Personally, I think no because the cost of keeping an attack goes exponentially high, so attacking any target, not only root website or whatever, for the first 30 minutes it might be doable but then second 30 minutes is much more expensive and then keeping it, for example, for a [a week], that's very hard and that gives a lot of time to operators in between to deploy whatever mitigation mechanisms they have. So, I don't think it's realistic that they can keep an attack so long that the cached zones are expired and then the distribution doesn't work, but that's my personal opinion.

DAVID CONRAD: And yeah, just to add to that, one of the points of hyperlocal is that you're moving the provision of root service closer to the source of the queries, so that in the case of like an ISP that is deployed hyperlocal, if there is an attack that's affecting the root service, which means the attack is against the resolver, that ISP

has the ability within their own infrastructure to track down and mitigate the sources of the attack one way or another.

That is sort of radically different to the situation we see at the root where we have about a thousand devices spread across the world that are not in a good – they do not know who their customer is to go to that customer and say, "Hey, stop beating the crap out of me."

BRAD VERD: Wes.

WES HARDAKER: No problem. I do have some concern over the focus so much on the root because the reality is, is that the issue is much, much bigger, and I think TLDs, for example, I think if they were taken out, might actually have more of an end user impact than the root service, so I guess I would encourage you to keep pushing back.

I suspect that the root kind of looks like a low-hanging fruit because it's more within reach, even though ICANN does have some arrangements with most TLDs that they could exert at least some help or pressure with, but I guess my real statement is I would really encourage you to remind them of that and then, also, your efforts toward dealing with protocol extensions and

other things that would actually help the DNS as a whole, means that the whole tree will stand up better rather than just the very tip.

DAVID CONRAD: So, two answers to that. One, the root is sort of the thing that ICANN.org has the greatest, because we run one of the root servers, sort of the greatest ability to influence, right? We can augment L root as we see fit, but the other answer is that the root is also unique, as far as I'm aware, in the DNS, in that in all other zones, the administration of this name servers for the zone are directly financially responsible, a responsibility of the owner of the zone.

In the case of the root server, that's not the case. The root server is operated by the name servers for the root are operated by 12 independent organizations and there is no mechanism by which ICANN the organization can have influence other than saying please, please, to change the way the root servers are operated.

We can ask politely, we can suggest, we can recommend, we can try throwing money at the problem, but the reality is that each of the independent organizations that are providing root service have to make the decision that's in their best business interest, so fundamentally the root is different in that way and it is, as far as I know, the only zone that that's through.

BRAD VERD: Oh, if I may. We understand that situation very clearly. We've been working on it for a number of years trying to resolve that. I think – and I'm going to try to channel Wes and some of the conversation we had earlier – I think what Wes was getting at was that while hyperlocal works for the root and for L, since ICANN can influence L, hyperlocal doesn't work necessarily for other TLDs that might come under attack with the same attack that you've laid out. Hyperlocal doesn't work for them to help mitigate anything and that is where it's quite possibly up quicker and larger impact, depending on the duration of the attack that was brought up earlier. So, Kaveh.

KAVEH RANJBAR: Thank you. So, two points. One is I agree with the statement but you have to keep in mind that root also attracts political interest, as well, so attack against root might have different angles than just being able to bring down part of Internet for commercial gain. There will be at least a lot of visibility into that, if someone does that. That's one.

And second, just to follow up on this presentation, after this was presented to the board, basically the board tasked Board Technical Committee to look into this report and ask for advice from rest of the community and come back to the board with

what they think. The BTC doesn't advise the board and this is not formal advice from at this point, they didn't request formal advice from, for example, RSSAC and SSAC, but they would really want to get feedback at least on this report or if there's more, through BTC from RSSAC, SSAC, and other possible interests of constituencies. That BTC hasn't yet formed a meeting, so I think next week it will happen and [inaudible] both RSSAC and SSAC will receive a request, and we will get a chance to basically mention all of these points. They will be all collected and presented to the Board, just [inaudible]. Thank you.

BRAD VERD: Thanks, Kaveh.

DAVID CONRAD: And just to reiterate something that Kaveh said. From the board's perspective, their view, again, channeling them, is that if there was an attack against the root that had sufficient impact to disrupt resolution on the part of the end users as a whole or even a significant portion where significant might be one person who happens to sit on Pennsylvania Avenue, then while some people might appreciate that, the problem is that the board believes, rightly or wrongly, that that would trigger concerns about ICANN's role in administering the stuff that ICANN does.

So, their interest is trying to protect ICANN's involvement in the operation of the Internet system of unique identifiers, of which coordinating root service is something that's specified in our bylaws, suggesting that we actually do have some role. If the congressional testimony occurs, it's most likely going to be someone from ICANN who is going to be the target of unhappiness, as opposed to individual root server operators, is the view of the board.

BRAD VERD: Thank you, David. So, that's actually a perfect segue into my question that I did write down during this thing. So, this was not pre-written – I came up as you were talking. So, in your words, you said a 10-minute outage results in congressional hearings and a bad day for ICANN and the board. We, RSSAC, have stated that if one goes down, it really has no effect on the DNS as the whole. That's a public document out there.

Kind of my question, really, is where along the spectrum of one to everything, is the trigger for a bad day for ICANN?

DAVID CONRAD: That's a really good question that I don't think anyone has an answer to. It isn't even… As we saw with DINE, DINE wasn't down for everybody.

BRAD VERD: If I maybe rephrase the question towards [inaudible] and maybe this is to Kaveh. Kaveh, is the expectation of the board that no letters are down?

KAVEH RANJBAR: Actually, that also came up during presentation. There was no clear answer but the board refers to root taken down, so what does that mean? And that was one of the questions that I asked – actually, what does that mean? The other thing, which was proposed, I don't know by which board member, but that maybe we should also look into what will have, assuming that it happens, so root goes down, whatever definition of that is, assuming it's all blackout. What will happen? So, what is the reality of the situation after that? And what people will do, how will it try to mitigate that?

Nobody knows the answers to that and that's actually something that we can work through our liaisons to the board to formulate the questions and then work on them in our constituencies, both SSAC and RSSAC, to basically try to figure out and then come up with what we think are the criteria, or, for example, what will happen, what is the scenario, assuming root goes down.

If we think that's a relevant question because some people in the board that maybe if you understand that, that helps us to better understand the actions, possible actions, or as someone suggested in RSSAC, maybe the board just needs to acknowledge that's a risk and we sign it off because the impact is very high but the probability is really, really low, so we are comfortable signing it off. But there is not enough information to assess this risk.

DAVID CONRAD: With that, I'm afraid I'm going to have to run. I actually have to be up at the CSG meeting. They're asking about the KSK rollover, so my colleague, Matt, may be able to answer any further questions. If not, feel free to drop me [a note].

BRAD VERD: Thank you for the attendance and thank you for fighting through the voice challenges.

TRIPTI SINHA: I think we've exhausted. Do you think we'd like some more time to discuss hyperlocal? We were going to delve into that and we've had that. Any other questions regarding that? Any other, if not, the other agenda item here says current planned research on the root service system by OCTO, so Matt.

MATT LARSON: I'm sorry, is too late? 30 seconds on the hyperlocal. The thing it wasn't completely obvious to me that the hyperlocal is not part of the mitigation. It allows you to survive the attack but it doesn't help mitigate the attack. Where one of the things we've thrown around that I didn't realize hadn't sunk in was my board's talked about get 10,000 Raspberry Pis and throw them around the world and somebody said, "Well, that's just like hyperlocal," but it isn't if you gave them all unicast addressing, so they would actually be creating a trillion tiny catchments that would make it hard for a DDoS attack.

UNIDENTIFIED MALE: Anycast.

MATT LARSON: What am I saying? Yes, anycast. It's been a long week. So, just that is a thought that's not hideously expensive and you would get well ahead of this thing rather than there's no way to win the arms race, that's a crazy way to win it. It's like going after an aircraft carrier with drones. It's different but it's worth thinking about.

TRIPTI SINHA: Thank you. Matt, do you want me to repeat what I said earlier? We're just curious to know what kind of research is OCTO looking to do on the root server system, how can we help you, how can we partner with you? Would you like data? Whatever.

MATT LARSON: Sure. Thank you. Well, first, let me say thank you very much for the RFC 8145 data that at this point I guess 12 out of the 13 letters are submitting. That's very, very helpful in analyzing the data for the root KSK roll. So, thank you very much for that. And looking at that, we're going to continue to look at that as we approach October.

I think in retrospect, we wish that were a different signal and what Warren and Jeff and [inaudible] talking about with sentinel I think is ultimately going to be a better way to measure because user impact is what we really care about, but we have the data that we have at this point, so we're going to continue looking at that.

If we can identify which of those addresses are ephemeral, belonging to VMs or containers, that would tend to indicate a lower impact because by definition, you can't have a bunch of users pointing at that as their resolver because it's got to be just the machine itself using the self-contained resolver.

So, if we're able to show that from the data, I think that would help alleviate at least some concern, I would hope, that we have such high values at the moment for what percentage of the 8145 capable resolvers are reporting only the old key, which hovers between now 20% and 25% higher than when we decided to postpone things back in last September.

So, that's one area of research that we're working on. In terms of things we've done recently, I don't know if some of you were at OARC last week and saw Paul Hoffman's presentation on TLDs occurring frequency that they occur at the root that was based on [inaudible] data, so thank you again for everyone for contributing the [inaudible] data as well as some more real time data, so that's an example of other work that we've done.

But I think maybe your question is what can the root operators do to help with our research? Is that…?

TRIPTI SINHA: What other questions are you looking to answer? Where's your research focus and is there some way we can partner? How can we help? Maybe there's synergies in what we're doing here.

MATT LARSON: Sure. As I said, the data is huge, so the statistics and the real time 8145 data, we're very grateful for, as well as everyone who

contributes to [inaudible], so I know I'm repeating myself, but those are very, very valuable.

In terms of specific other requests, I can't think of anything at the moment, but I can certainly take that back to the research team and relay the request, the interest from RSSAC and possible partnering or helping and see what I come back with.

TRIPTI SINHA: Does anyone else have any questions of OCTO? Or do you have any questions of us?

MATT LARSON: I don't have any questions.

TRIPTI SINHA: Well, let me check the Adobe [inaudible].

WES HARDAKER: I have a quick one or just a statement. Matt, at USC, as we have a research project underway to develop some [BS] tools for protecting DNS critical infrastructure, and we are probably four to six months away from releasing some code for some of those, including some whitelisting type patterns that I think have never been looked at before, and we'd be happy to turn those over to

you guys, if you guys want to play with them either combination with [inaudible] or just by yourself at some point. This is a FYI.

MATT LARSON: Okay. Thanks. I'd be interested in that. I mean, ultimately, that would be the [inaudible] ICANN's DNS engineering team that would decide whether or not to use those in production, but thank you. Yeah. Please let me know.

TRIPTI SINHA: I guess we're done. We adjourn early today. Thank you very much.

MATT LARSON: Few people complain about meetings ending early, so thank you.

TRIPTI SINHA: All right. Thank you. Meeting is adjourned.

BRAD VERD: Thank you all.

[END OF TRANSCRIPTION]