ABU DHABI – Joint Meeting: RSSAC and OCTO
Tuesday, October 31, 2017 – 17:00 to 18:30 GST
ICANN60 | Abu Dhabi, United Arab Emirates

BRAD VERD: Sure. Hello, everyone. This is the joint meeting of RSSAC and OCTO, Office of the CTO. Are we pulling up the… I know we have a few questions between the two. So, number one on the agenda is root zone KSK rollover, and we'll talk about RFC 8145 key tag data and Any Other Business. So, before we go any further, is there any other business that we should add right now to the agenda?

All right. Hearing none, we will start with the root zone KSK rollover postponement, and I'll turn it over to you, David.

DAVID CONRAD: Thank you. So, we postponed the KSK rollover. Next slide. Done. It's been a long day.

Okay. So, as I'm sure all of you are by now aware, we had decided on September 27 to postpone the KSK rollover because we were seeing data that we did not fully understand. We decided to postpone in order for us to get a greater understanding of what was actually going on that was causing

the higher-than-expected percentage of KSK 2010 only announcements coming in through the RFC 8145 new feature that was apparently deployed by default on [Unbound], some version number, and by default not on Unbound some version number.

Since that time, Unbound has turned on that feature by default, and we're beginning to see more data that suggests that the percentage that we were seeing back at the end of September is actually going up. Duane, du you want to talk about this a bit, because it's a lot of your data?

DUANE WESSELS: Yeah, sure. Thanks, David. So, you're right. Looking at the data that has come in since October 10, October 11, there is an uptick in signals showing only the old KSK. I can't say confidently, but preliminarily it seems like that is due to reports coming in from Unbound software that we weren't getting before. I don't know exactly why. These additional reports from Unbound seem to have a lot of sources stuck on the old key, but we're investigating that. Roy and I are collaborating and looking into that. That's about it.

There was one other curious thing in that data, which I'll talk more about at the DNSSEC workshop tomorrow if anyone wants to go there and hear that story.

DAVID CONRAD: Cool. So, at this stage, we're still in data-collection mode. We have hired one of the usual suspects, some crazy person named Joe [Abaly] and his company to help us. We take the original 500 IP addresses that we had seen that were announcing the KSK 2010 only. We hired Joe and his company to try and track down the resolver operators and understand why the servers were reporting KSK 2010 only.

The preliminary information is along the lines of what we would expect, where sometimes the resolvers were operating on a read-only file system, so it couldn't write out the key tag. And I gather it was obviously not a fatal error for that to occur. A number of folks had used the trusted-key statement and bind instead of manage keys, and must actually appear not to know the difference between the two. If I remember correctly, the preliminary data was suggesting that about half the folks had not responded to the requests. So, that's an ongoing effort.

At this stage, we're still in data-collection mode, and we're not even guessing when we will be able to re-initiate or continue with the rollover. Obviously, we're going to be waiting at least until the next quarter, just because of the physics of the way the key rollover ceremonies work. But, I will be quite surprised if we

have sufficient data at that point to be able to make any sort of determination about how to move forward.

It actually opens up an interesting question about how we will move forward. The numbers that we're getting are obviously resolver numbers. The design team and the operating plan had indicated a 0.5% of end users that would be negative impacted. And I'm unclear as to how we will translate the number of resolvers into the number of end users. So, there's probably some additional discussion that we're going to have to undertake. How we do that is still up in the air right now.

There've been a number of suggestions from just going ahead and breaking those folks who aren't updating the key correctly to going out and doing a multi-decade public PDP or some such crap like that. I don't know. I am actually joking about that latter one.

But that's something that we will undoubtedly be consulting with RSSAC and other with about how to move forward, based on the data that we're receiving. So, I guess the short answer is stay tuned, because we're continuing to try to understand what's going on.

Any questions, comments, screams of outrage? Speaking of screams of outrage, Roy, welcome. We were just talking about

the KSK rollover. Is there anything that you would like to say, since you're the one who is providing a good chunk of data?

[ROY ARENDS]: I'm not sure if this was mentioned yet. Duane has written an extension to his RSSAC KSK – maybe you're going to talk about it later.

DUANE WESSELS: I assume that was the next topic on the agenda, but I'm not sure.

[ROY ARENDS]: Great, and that's what I'd like to point out. Go ahead.

BRAD VERD: Before we get there, the first question I have is "Is there anything further you need from the root operators that we should take back and get for you?" Second question "Given the delay," I've been asked a couple different venues – not venues. Different areas that "How long?" And I know this is a bit of a loaded question, but "If this is extended for an unextended period of time, are we going to run with the two keys with the packet size the way it is? When do we make a decision on that?" That type of thing.

DAVID CONRAD: So, answering your second question first, my understanding, which acknowledged experts can correct me if I'm wrong, is that since we have deployed with a larger response size and so far as I can tell the Internet does continue to work, and we've actually not received any indications of problems with the larger response size, that we can continue in that mode as long as we need to. At this stage, I don't believe there is any justification for going back to, what was the state, D? C? Whatever. Yeah, going back to C. I don't see there's any need, any reason to do that.

UNIDENTIFIED MALE: I think we're on C now.

UNIDENTIFIED MALE: No, we're on D.

DAVID CONRAD: Yeah we're on D now. Based on that, my assumption is that we can continue on this mode as long as it's necessarily for us to get a reasonable level of confidence about moving forward.

UNIDENTIFIED MALE: Thank you.

DAVID CONRAD: And I have, of course, forgotten the first part of your question.

UNIDENTIFIED MALE: Is there any further you need?

DAVID CONRAD: Oh yeah. Thank you. Yes, actually there is. We are continuing to need data, particularly the key tag data. So that is actually, I believe, a nice segue into the next question.

DUANE WESSELS: I think we've already talked about this a little bit within RSSAC already, but number two on the topic list is a request for key tag data, not just from those that are providing it now, but from almost all the operators that we can get it from. This takes the form of the DNS cap plug-in called RZ Keychange, which a lot of the operators are already deploying.

I've done some work to modify this so that it was also report the key tag signals to ICANN in this case. I think there is still an action item on maybe Roy and I to finalize that, get it polished up and then send out some e-mails and ask people to deploy that and start providing the data.

I know [Terry's] looking at me funny, because he doesn't like to run this plug-in, because it uses pcap files. Can you remind me,

is ICANN running this plug-in already? Are you doing something different?

TERRY MANDERSON: We are not running the plug-in. We are providing every query response, etc., to OCTO in CBOR and they have everything as far as we know.

DUANE WESSELS: So, you can get that data from them that way, I guess.

UNIDENTIFIED MALE: The team from [inaudible] were kind enough to implement very, very similar if not the same of your previous RSSAC KSK – the number of keys squared, etc., etc. So, we have that exact same information from L-root that we have from all the other roots. The key tech data, I have a slightly different process for that, but that's also –

UNIDENTIFIED MALE: Okay

DAVID CONRAD: Okay, great. Sorry, go ahead.

WES HARDAKER: Thanks. Wes Hardaker, USC. Right now the data, and I'm very happy to provide it and I'm very happy to run the next version that gives you more information. The one issue is it's sort of a black hole at the moment from our perspective. We give you data, and we haven't seen much in the way of results or analysis or something.

In the previous version when we did this, we were able to make sure that data was actually being captured. One of the things I realized last week, I'm like "Hey, I wonder if they're still getting our data that we're sending them?"

We have no ability to actually see. I actually have no idea if we're dropping off and you just figured we turned it off or something like that, whereas previously we could at least see graphs that showed that we were still spinning data.

UNIDENTIFIED MALE: Currently, the data shared between the root management partners, [Verisign] and ICANN, actually, I think OCTO and [Duane]. But I'm happy to provide anyone. The only problem is I can give you your data, but I'm not sure that others would like to give you their data. I have no idea how it works.

WES HARDAKER: I think we can come up with a joint statement that may relieve you of that responsibility.

UNIDENTIFIED MALE: I'm happy to share. What I was referring to is very nice graphs.

WES HARDAKER: I like nice graphs.

DAVID CONRAD: Yeah, one thing. I should have mentioned, the previous item is that we're right now wanting to get in contact with the other resolver vendors to see if we can encourage them to develop the key tag data stuff so that we can get… One of my big concerns right now is we only have a subset, and actually a subset of a subset, of resolvers that are actually configured for DNSSEC that are announcing this. So I want to try to get other resolver vendors to deploy code that does this as well. Obviously, there's going to be the normal deployment long tail on this, but we've already been surprised that actually people are deploying this stuff. So maybe we can get surprised again.

BRAD VERD: So, is there a formal request from OCTO to make this change, or is the informal request good enough and we'll manage through this and you guys will just start getting your updates?

DAVID CONRAD: I think we did a formal request a while back, about just getting data for the KSK roll. We can just say that's still in force. Until we actually roll the key, any data that you all can provide is something that is undoubtedly going to be helpful.

BRAD VERD: All right, Wes.

WES HARDAKER: I assume someone will let us know when the code is available.

DAVID CONRAD: Yeah, for sure.

BRAD VERD: All right. Anything further on item number two, RFC key tag data? All right, moving on number three. Any research on the root server system? This is a question from RSSAC to OCTO. This is basically is there any research or something going on that we're not aware of that we just –

DAVID CONRAD: We do a number of sort of ad hoc research projects. For example, as a result of ongoing discussions with regard to [corpo and mail], Roy had scrounged root-query data looking for information related to [corpo and mail] to aid in discussions at the Board level about the number of queries and whether things had changed in any significant way from 2013 when the [Interall Study] was looking at name collisions. So, that's the continuing thing. The resulting paper actually will be published momentarily. It kind of fell through the cracks. We also sort of root server system related, but not really. We're doing analysis regarding two-letter country codes at the second level.

So, we do a bunch of these ad hoc things based on requests from either the Board or from specific requests from community members. We don't currently have any specific efforts underway regarding the system, architecturally as a whole. One of the items that we do have on our research agenda is to try to develop a methodology to identify good locations for the deployment of instances, because ICANN has a lot of instances. We will be working with Terry and his team on that project moving forward. Other than that, I'm not thinking of anything. Am I missing?

UNIDENTIFIED MALE: We've got a few ideas that haven't crystallized out yet, so there's nothing to see there yet. I think that's it.

DAVID CONRAD: Any time any of these projects get more formalized, we'll let you all know. I apologize for not letting you know about the [corpo and mail]. If anyone cares about it, I can send you the draft paper. It's not particularly exciting. I suspect you all know what the answers are there.

BRAD VERD: Great. Please share and we'll forward it along. Any questions regarding that?

All right, moving on, number four. Any communications with other organizations about the root server system, including placement?

DAVID CONRAD: Oh yeah. Yeah. Yes.

BRAD VERD: Can you elaborate, just for everybody's benefit?

DAVID CONRAD: So, I recently did a whirlwind tour of Asia, visited a number of countries. In most of those countries, I did get questions about root server availability and placement that continues to be a topic of interest. I actually believe – I think because we mentioned in the Board RSSAC session, we are seeing a lot of interest with regards to the root server structures, root server infrastructure, architecture, how to play in the root server system, how to become a root server, all those sorts of questions. None of these are particularly new, but I believe – a personal opinion, no data to back it, so with those caveats – I believe there is an increased interest in the system as a whole as a function of the increased and, in my view, quite impressive amount of work that RSSAC has been doing of late. The various documents that have been coming out, the various statements that have been made, all of that I believe has been extremely positive, although it does have the side effect that it does bring attention to yourselves. And, when that happens, then you end up getting questions. And this is not a bad thing. Actually, I think it's a very, very good thing, because it's beginning to de-mystify a lot of some really amazingly bad assumptions about how things work.

So, I would encourage people to continue to be open and document processes and architectures and structures and relationships and all that sort of stuff.

I will say that in one of the meetings that I had, it started out with about a 45-minute monologue on why that particular country deserved a root server, and then ended with "But, if all the letters – if it just became rootservers.net and all the letters no longer existed, then we wouldn't have any questions," which I thought was actually a really interesting way to end that discussion given that I'd just sat through a 45-minute discussion of how important it was that that particular country had a root server.

So I applaud, in the greatest possible terms, RSSAC's movement to at least in public discussions move away from declaring the letters to be more of a cloud service.

UNIDENTIFIED MALE: If possible, are you able to communicate those really bad assumptions to us, so that we might be able to address them in some of our documentation?

DAVID CONRAD: Sure. I will be happy to. I will actually draft up a summary of some of the discussions that I've had and send it to RSSAC.

UNIDENTIFIED MALE: Thank you for that. [inaudible]. In your travels and dialogues, do you ever mention RFC 7706?

DAVID CONRAD: Every time.

UNIDENTIFIED MALE: Okay, and is there any effect in your mentioning it?

DAVID CONRAD: It depends on the context. So, in the cases where I've been speaking with governments, it has absolutely no impact at all, because those discussions are not technical. In the discussions that I've had in the context of network operators and network service providers, a number of them have already known about it, but they weren't entirely confident about what it meant and all of those implications.

Actually, thank you for asking that question. One of the issues that they have raised in those discussions has been the robustness of the zone transfer infrastructure, because if they're going to be deploying a mirroring of the roots and they want to have the ability to gain the root zone, and the root zone transfer infrastructure right now is laid out as a little less than well

documented. You know, ftp.internet.net is not usually good answer to "How do I get the zone?"

So that's something that at some point it may be worthwhile to explore, particularly as, which I assume will be the case, 7706-like mechanisms get more and more deployed. Building out the zone transfer infrastructure may be something that would be worthwhile to explore.

BRAD VERD: All right. Any further questions? Wes?

WES HARDAKER: Two things. One, this was actually not my original comment. So, with respect to 7706, I've recently completed standing up a 7706 helping service that I call LocalRoot. It's called localroot.isi.edu and people can go sign up and they'll get a TSIG key and transferability and things like that. It walks them through at least the bind configuration currently that they need to actually to 7706, and how to put it in a resolver. That's just FYI. I'll be talking about it tomorrow at the DNSSEC workshop. We'll be slowly announcing it wider as time goes on.

The other thing is, in your travels and conversations, I think one of the biggest misunderstandings that we consistently run into is that having – there's a lot of discussion around people needing

something in their region, and they need one thing in their region. From a technical perspective, I'm sure we all know that "No, you need three to four to five instances in a region, from different address points." That's really the only way to get proper service. There's been more and more studies with Anycast, some done by my colleagues at USC that have shown that there's actually some magic braking point of how many you need.

I just reviewed another paper in the last month that talked about actually how much benefit you get beyond so many instances in an area. Is there any way to convey that need to them so that they don't just talk about needing to control one for themselves, but rather that they actually need a collection in their area in order to get robust service?

DAVID CONRAD: That topic does come up. In the more technical discussions, there is an acknowledgement that they need more than one. They don't know how many. So, if you can provide a reference to me or that paper or those papers or that research, that would be really useful for those folks.

In the less-technical discussions, as I'm sure you're aware, the discuss usually boils down to "I need a letter because I'm important." Right? And it actually has nothing to do with the

robustness of the infrastructure, it has to do with the fact that they're important, therefore they need a letter.

So, in those discussions, I usually try to leave.

WES HARDAKER: I'd be happy to provide you references to a couple of papers that I know that would be interesting.

BRAD VERD: All right. Any further questions? All right, we're on to Any Other Business, which there was nothing listed. David?

DAVID CONRAD: I did have one, and I think I raised this at the last meeting. I don't recall, because my brain is tapioca right now. I'm forgetting the words now, because my brain is tapioca.

Server ID, that's it. Server ID information. My understanding is that the majority of root servers do provide server ID information to help identify instances from out-of-band instance identification; however, not all do, if that's right. And most root servers also do the in-band EDNS0 option to provide identification of instances. But again, not all do. I'm not asking about – it's sort of policy decisions about yes or no, that kind of thing. I'm mostly interested in what the operational rationale is

for not providing the server ID information. If there is an operational consideration like there's a potential for security implications or there's some other potential that could operationally impact the other root servers, it would seems that that would be useful information to have. If it is purely a policy decision by higher-ups of one form or another, that's fine, too. I'm not suggesting that that be changed. Just I'm interested in understanding if there is a security stability potential that some folks believe that others may not, that that would be an interesting bit of information to bring to light.

BRAD VERD: I think that's a reasonable question. And I don't want to have this sound like I'm punting, but I think it would be great if you could put that in a formal request that you could hand to RSSAC. And since it is an operational question, we will take it to ROOTops, and we will try to track down an answer for you.

DAVID CONRAD: Great. I will ensure that someone reminds me to actually put that into a formal thing. Cathy?

BRAD VERD: Anything else on that topic from anybody?

DAVID CONRAD: I have another Any Other Business that may be of interest to you all. So, my esteemed colleague to my right will be speaking about the frequency of key roll at the DNS workshop tomorrow. Ignoring the fact that we had to postpone the existing key roll, eventually it'll happen, we assume.

But right now, there is some ambiguity about the frequency. The existing DPS, I believe says, well I know it says that we will roll the key after five years, or some miracle happens and people are able to factor the key, or we lose it or whatever. But there is some question about whether that means we roll every five years or whether we roll maybe not again. Or maybe we roll every 15 minutes, because I know that's what Dr. Crocker actually wants.

Roy will be talking about that at the DNSSEC workshop tomorrow. If you have interest in that topic, please hunt him down and [laugh at him].

BRAD VERD: All right. Any questions? Wes.

WES HARDAKER: I keep monopolizing the microphone. I'm sorry. One of the – and we discussed this in some other meeting the other day, too, but one of the other questions in my mind is, going forward to the next key roll, will there be a re-evaluation of the current key rolling procedures as to whether it's adequate given our lesson learned with this one. I wasn't going to bug you about it today, because I figured "Let's get over the first one before we start designing the next one."

But, since you asked, 5011 has shown to be interesting in order to actually do a key roll. The way the current design for rolling is done, it's actually not even done with the guidance with the original author of 5011. There's probably better ways to go about doing it, and I think it might be worst reconsidering how we go about adding another key and then when to remove it. They may not be done in the same time zone.

DAVID CONRAD: Yeah, so we have learned a number of lessons. I'm hoping that by the time of the next key roll, we'll have a better measurement infrastructure available to us so that we won't have these sorts of surprises that we had in the past. We have had some very, very preliminary discussions about the need to have a standby key, which would imply having an additional KMF and all sorts of fun stuff along that line.

There's clearly going to be additional discussions. At this stage, I'm sort of seeing this KSK rollover as more of almost an experiment, because everything was new. We had very little understanding about the potential implications and all of that. "Was the response size going to be a problem?" Right? So, those sorts of questions.

So, we've learned quite a bit as we've been going through this, learning how to change the engines of the airplane as the airplane is in flight. But it may be that the next key roll may actually also be an algorithm roll, because maybe by the time we'll have ECC algorithms that are deployed sufficiently to actually make it useful.

There are a lot of questions that need to be answered. There's a lot of potentially additional infrastructure we need to explore. There are so many ways that we can think of to break the Internet that I look at the future with glee.

BRAD VERD: Please share your glee with everyone else.

Any Other Business? Any other, other, other business? David?

DAVID CONRAD: Beer time? No. I don't have any other business. Anyone else have any other business?

BRAD VERD: All right. With that, we'll adjourn. Have a wonderful evening.

DAVID CONRAD: Yeah. Thank you very much, and enjoy the Steve Crocker tribute.

[END OF TRANSCRIPTION]