
OZAN SAHIN: Good morning, good afternoon, and good evening. This is the RSSAC023 History of the Root Server System Update call held on the 6th of January 2020 at 16:00 UTC.

On the call today, we have Anupam Agrawal, Brad Verd, Chris Ishisoko, Dessalegn Yehuala, Duane Wessels, Fred Baker, Wes Hardaker, Hiro Hotta, Ihtisham Khalid, Jeff Osborn, Karl Reuss, Kazunori Fujiwara, Kevin Wright, Mukund Sivaraman, Paul Hoffman, Paul Muchene, Ray Bellis, Russ Mundy, Ryan Stephenson, Shailesh Gupta, Shinta Sato, Michael Frauenhoffer.

From staff, we have Danielle Rutherford, Andrew McConachie; and myself, Ozan Sahin.

I'd like to remind you all to please state your names before speaking for transcription purposes. Thank you and over to you, Paul.

PAUL HOFFMAN: Great. Thank you. So, Ozan has sent out a couple of messages about the call and they have links to the main document and Andrew posted a link to the document this morning.

Basically, today's call, we will go through the document, catching just the high points of the comments that are there. Our goal is to close out the document as soon as we can. That is, we're not trying to keep this open super long. Basically, what we're going to want to do is try to get everything we can reasonably complete, do sort of a last call, and then we, being the caucus, will hand it over to RSSAC for finishing.

Note: The following is the output resulting from transcribing an audio file into a word/text document. Although the transcription is largely accurate, in some cases may be incomplete or inaccurate due to inaudible passages and grammatical corrections. It is posted as an aid to the original audio file, but should not be treated as an authoritative record.

So, Ozan, if you can—ah, very good. Thank you. There we go. We have the document up. So, does anyone have any particular issues before we start or questions on the document?

Okay, seeing no hands and such, we can go ahead. So, the first overarching part is that, in section three, there is the history of each root server operator. So, Danielle, can you say a little bit about the status of that section?

DANIELLE RUTHERFORD: Yeah. Hi, everyone. I sent out a few reminders to everybody to update your specific root server section in chapter three. I have not heard back anything from Cogent or Verisign on if those operators want to change anything to their sections. I'm still waiting on ISC. I believe they asked for a couple of weeks in December but I haven't seen any new edits to that section yet. And then everyone else has either provided edits or confirmed that they link their section as-is.

JEFF OSBORN: Actually, the changes were fairly minor but they've been in there for a couple of weeks.

DANIELLE RUTHERFORD: Oh. Sorry about that, Jeff. Okay.

JEFF OSBORN: No problem.

DANIELLE RUTHERFORD: I recorded that.

JEFF OSBORN: We had a huge amount of effort and then wrote a bunch of changes and then really minimized them after a lot of thought. I apologize I didn't flag it.

PAUL HOFFMAN: Okay, great. So, Danielle, you said Verisign and Cogent were the two who you had not heard from either saying everything was okay or that they wanted changes.

DANIELLE RUTHERFORD: Correct.

PAUL HOFFMAN: Okay. So, I don't see anyone from Cogent on the call. I do see Brad. Brad, can you sort of take charge for Verisign there?

Well, you're still on mute. But I will say, Brad, please get any updates you have to Danielle. And we'll send out another message to the list on this.

Okay. So, going back through, Ozan, if you can scroll up to the table of contents. What we have—scroll down a bit right there. Section three.

Terry Manderson had an issue with the names of the subsections in section three, that there was some discussion in RSSAC earlier that never really came to any fruition about not using letters, using letters, using alternate names and such like that. Then there's a little bit of back-and-forth in the document, in the comments, starting with Terry's.

We still are living in a world where people think of the root server operators by their letter names, plus maybe an additional name—or at least the organization's name. So, the question really is what do we want to title the sections here that will appear, of course, both in the main body and then in the table of contents.

Given that there was really no resolution in RSSAC—we never got an RSSAC document talking about the naming and such—it may be best just to let each individual operator decide what they want as their title. I'm not a big fan of all capital letters. So, for Terry's change here for L.rootservers.net, that's just a personal one. I would certainly make that lower case if you wanted it to be domain name.

But anymore discussion on this? Are there other RSOs who would want a different name or anything like that? I'm not seeing hands.

WES HARDAKER:

I don't think I can raise a hand but that's a hard one. The problem is that we are talking about the history over time, right? So, the names that are there right now might be the current names, or depending on how you look at it, the [just passed] names. Even early in the document, the names changed over time. So, that is listed as current root server operators.

I don't think there's really any documents that say A—root in them, whether that be RFCs or RSSAC documents. No, there might be some RSSAC documents.

I think that L.rootservers.net is probably the authoritative current name. So, I think Terry's change makes sense. I agree. I like lower case better myself. But the current names as identified in the root zone are L.rootservers.net. I guess I'm suddenly leaning toward that which I was not at the beginning of my sentence.

PAUL HOFFMAN:

Actually, what I proposed was that each root server operator gets to choose—

WES HARDAKER:

Yeah. Let me push back on that then, right? I think consistency in this document would be better than letting everybody choose their own thing. And the reality is each of the service identifiers that are currently in the root zone, which this document is sort of talking about, are already named in a particular way and we ought to stick with that if we're talking about current.

PAUL HOFFMAN:

Fred? Fred, you had your hand up.

FRED BAKER:

Here we go. So, there have been discussion in the RSSAC about moving away from the letters. I think that's kind of the general consensus. But you're right, we haven't written a document or given direction in that regard.

If it were my call, which it's not—it's a group decision—but if it were my call, I think I would replace all of these with the name of their company (mumble.rootserver.net). So, that would be Verisign and USC and Cogent and so on.

The reason being that the term A-root identifies the root server, if you're part of the club and you kind of understand what the letters mean. If you're somebody else, somebody that's actually likely to be reading the history document, saying that it's A, B, C, D, E or whatever, probably doesn't communicate much information.

So, it seems to me like the best bet here is to think about the reader and put it in something that's going to be interesting and useful to the reader. So, that first one would read, for example, Verisign (a-root).

PAUL HOFFMAN:

Okay. Russ?

RUSS MUNDY:

Yeah. Thanks. One of the concerns about the approach we're using that approach that Fred just discussed is that some of those names have changed over time. We can list who currently is the operator but some of them have changed organizations over time—and that's noted down in the text.

My inclination would be to perhaps do a concatenation of what was I think Wes's and Fred's suggestion and that is to do—it's probably a relatively long title but the name from the authoritative root zone and the organization because section three says current root server operators, which does, at least in my view, reflect the organizations make more sense but to ...

I think a broad swath of certainly the technical part of the Internet and even those beyond think of it in terms of the letters. So, perhaps having both the authoritative zone name and the current organization name would be a way to handle a compromise here.

FRED BAKER:

In that case, Russ, you're agreeing with me. You're saying that it would have the organization name—Verisign, USC, whatever—and then concatenated with that I guess it would actually be the domain name as Terry wrote his.

RUSS MUNDY:

Yeah.

FRED BAKER:

Yeah. So, a.rootervers.net.

RUSS MUNDY:

Right. Yeah.

FRED BAKER: Yeah. And I was suggesting put that in parentheses just to make the difference be obvious.

PAUL HOFFMAN: Duane?

DUANE WESSELS: Yeah. Hi, this is Duane. I just wanted to express support for what Fred was proposing. I like that.

PAUL HOFFMAN: Great. And doing it that way also makes the document a little bit more searchable, so that if the reader—as Fred was saying, we should be thinking about the reader here. If the reader knows the name of the organization and they want to see the history, they might search for that. Or if the reader doesn't know names of organizations but had heard something about Q-root or whatever, that they could search for that as well.

Now, they won't be able to search for Q-root but they would be able to search for q.rootserver and then have a good chance of finding it.

WES HARDAKER: Just for reference, I like that, too.

PAUL HOFFMAN: Great. Ryan?

RYAN STEPHENSON: So, are we going to change the name? And I apologize if I'm kind of confused on terms. They're kind of root server operators but they're also root server ... If we go with the L.root-servers.net, will we change that to root server indicators then? Or identify ... I'm sorry, not indicators—identifiers—versus root server operators?

PAUL HOFFMAN: So, the section is about root server operators. Root server operators can be identified in many ways. One is by their organizational name. Another is by a domain name that seems to be common and such like that. So, we could keep the section title the same because what this is, is the history of the operators.

RYAN STEPHENSON: Okay, awesome. Thanks.

PAUL HOFFMAN: Other thoughts on this before we move on? And I will start a separate thread on the mailing list about this since people can ... Naming is hard. That's why we have a DNS, so people will tend to want to discuss this further. Andrew, I see your hand up.

ANDREW MCCONACHIE: Yeah. Sure. The only thing I want to mention is that these terms are used throughout the document as well. I guess this discussion has been constrained to just the table of contents.

PAUL HOFFMAN: Plus the section headings they're derived from, but yeah. So, when you talk about other ways of identifying an organization or a root server, we may have to look at those as well.

But, for this, when this is really the subsection that is about them to make the naming clearer there.

ANDREW MCCONACHIE: Okay. That's clear.

PAUL HOFFMAN: Great. Then let's move on. So, Ozan, if you can scroll down to the end of section one ... A little bit higher than that, sorry. So, this discussion here—wait, you just missed it. Keep going. So, it's the highlighted section here.

There was a question earlier about do we want to ... Should this document also discuss the root zone as compared to the history of the root zone operation?

So, there were things such as the contents of the root zone, when it was signed, the introduction of IDNs, the greater size and such. That can happen here, although that is quite different than the history of the

root server operators and operations, simply because we believe and we state in this document that in fact the root server operators have been quite consistent about always serving the root zone. So it doesn't really matter what's in it.

One suggestion I heard, I think it was at the last caucus meeting, was that in fact maybe RZERC wants to start a document about the contents because they're the ones who nominally oversee changes in the contents.

But this suggestion has come up a few times, so I want to talk about it here. Does anyone have any strong feelings one way or another on should this document also cover contents? Because if it doesn't cover contents, that's easy to describe in a sentence or two saying it doesn't cover contents. It's a little bit funny doing that in the absence of a document that does cover the history of the contents, but the history of the contents will certainly add a lot of time to creating this document. So, thoughts?

FRED BAKER:

It might be good to talk with the chair of RZERC and just see if they have any plan to produce such a document.

PAUL HOFFMAN:

Right. Yeah. That sounds good. And even if they don't, someone else might produce a separate document. I don't know if that would be RSSAC caucus or an individual. This certainly seems like a reasonable academic exercise as well. Paul Muchene?

PAUL MUCHENE: So, what if you included the content or mention of the root zone in the appendix, just of this document, as a way to perhaps a more elaborate document that will give the contents and the history?

PAUL HOFFMAN: So, that's a possibility, although then, as you can see from the highlighted paragraph here in yellow, what people consider to be important in the contents is different from person to person. I happen to be an IDN person. I consider that more important than adding IPv6. I can see Fred's virtual hand raising quickly to probably disagree with me on that.

So, it would be challenging for us to limit that, although we could list many kinds of things. That's one thought. Russ?

RUSS MUNDY: Thanks, Paul. I would be in favor of the short description that says this document does not address the contents of the root zone or the history of the contents of the root zone simply because it is a very challenging effort to try to do ... If you're going to do something like that, you should do it thoroughly and there's a lot of work involved, so I think we are much better off to just say deferred for later work or other activity.

PAUL HOFFMAN: Great. Duane?

DUANE WESSELS: So, as chair of RZERC, I'm happy to bring this—

PAUL HOFFMAN: Thank you. I am removing the note that I wrote down for myself about talk to the chair of RZERC.

DUANE WESSELS: Yeah. I will bring this up to the committee, to RZERC. However, based on my experience with RZERC so far, this seems like ... It's unlikely to me that they would want to take this on, just because in general RZERC has been very, very careful about the work that it does and what it produces. I think most people would think this is out of charter, out of scope. But I will certainly bring it up and see what people say.

PAUL HOFFMAN: Thanks. So, here's a general question to the group. We have a short list of interesting things that have been added to the root zone over time here in yellow. There are probably some other things people thought of. Are people aware of any documents out there now, either maybe in the IETF or individual research that's been published academically, tracing the content by anything other than just the size of the root zone? I know there's been some studies of the size of the root zone. But the types of records, the introductions of new types of records, such like that. Is anyone aware of these?

Seeing no hands. I'll bring that to the list as well because if we are going to punt, it would be nice to punt at least with a reference to a forward section somewhere else if we can. But this will also be a topic for the mailing list.

Okay. Ozan, if you can scroll down to section 2.8. Oh, and we're having a numbering issue here, but yes, that would be the section 2.8 I was discussing is the new addition.

So, this was discussed earlier and it was reasonably controversial. I don't think we are done discussing it yet, but I would very much like people to take a look at this short section. I'm sorry, people who were around in 1998, so firsthand experiencers which does not include me. And see whether this current wording is acceptable to them. If not, to propose changes.

At some point in the near future, we will start an editing pass to reduce the number of edits. That is, we'll accept all of the obvious edits, so that the non-obvious edits in the comments stand out a bit more.

But this particular section was discussed and I know there are strong opinions about it. Again, my preference—I mean, anyone can comment of course—would be that the comments come from people who were either root server operators or root observers at the time of 1998 to talk about this, because it is more of a historical ... Again, our document is meant to be historical, not opinionated about what should have happened or what might have happened. So, please take a look at that.

Then, Ozan, if you can ... I'm sorry, were there any comments on that at the moment?

Okay. Then if you can scroll down to 2.9 or the new 2.9 on the IANA transition. This is also a new section that I think can use fresh eyes. So, same thing to the group. I think most of us have been here since about 2014, that there could be more views on this. Please take a look at this.

Specifically, are we describing factually what happened? There were certainly a lot of feelings about the IANA transition. And I will nudge people in IANA, especially Kim Davies, to take a look at this as well. But this is another new section that we want to have good eyes on. Any comments on this one that people have thought of at the moment?

Okay. Seeing none.

UNIDENTIFIED MALE: There's a hand from Russ.

PAUL HOFFMAN: Oh, I'm sorry. Russ, please go ahead.

RUSS MUNDY: Hi, Paul. I think that there does need to be maybe some wording adjustments here and it would be probably good to look at it also from the perspective of some of the wording that was contained in the [ISCG's] plan that was sent to DOC that could perhaps get better realigned. That these words could be in a little closer alignment with what was said that was forward as the results in I think it was the March date that was eventually a plan that was accepted.

PAUL HOFFMAN: Okay, great. Can you put a note here about that? And then also say whether you want to replace the words here with that or augment them.

RUSS MUNDY: I'll put a note. And I think it's probably a little bit of both. Kind of subtleties that some people might take exception to from some of the current wordings.

PAUL HOFFMAN: Okay. Very good. Thanks. Other thoughts on the IANA transition wording? Okay.

Then the last bit was that there was a proposal—I think it was from Brad—that we add a short section on the 037-038 thing, which is listed here as 210. And by the way, we will do a renumbering before too long, so don't get too attached to these numbers because we now have multiples for a couple of them.

However, there is no wording here. This was a proposal. Would anyone on the call like to take a stab at writing the wording for this? Now is a good time to volunteer. Okay, I'm not going to do the IETF thing of assigning a writer, but please consider, especially those of you who have been active in RSSAC versus RSSAC Caucus on the development of 037-038.

We're not looking for a lot. I think that a paragraph or possibly two would be good, and especially because it will be dangling. At this point, we would love to get this document out sooner rather than later, well before the GWG has finished its work. So, Andrew?

ANDREW MCCONACHIE: Staff can take this action item if the work party is okay with that.

PAUL HOFFMAN: I would be fine with staff doing that. I sort of would have liked to see work party activity on this, but I didn't. So, unless someone put something in the document in the next, say, week ... Yeah. Staff can you take a swing at that?

And again, it can be shortish and simply say watch for further developments.

WES HARDAKER: Actually, I was going to suggest ... I think there's text that might actually fill that out of 37 and 38 and I'd be very tempted to quote as much of that as possible and just pull out a paragraph saying this is what it says about it.

PAUL HOFFMAN: Okay, great.

WES HARDAKER: I don't know if anything needs to be added, but my belief is probably not.

PAUL HOFFMAN: Okay. Well, I would say at least one thing would need to be added, which is that the GWG currently is already being formed. So, it's not just 037 and 038, it's also the action plan.

WES HARDAKER: Right. That's fair.

PAUL HOFFMAN: Good. And then, Fred, you sent me a private message on the process question but I can say it here. Fred asked how do I see this document developing. I believe that we still have open issues and I will prod folks on the mailing list about that. And I would hope that we could be done-ish in maybe a month depending on hearing from Cogent, hearing from Verisign, and activity on the things that we just covered, the open wording.

So, I'm not looking at a hard deadline on this because the document has existed for many years without an urgent need to update it, but I certainly don't want it to linger.

My picture is that we hopefully will be done before the Cancun meeting, so maybe it can either be discussed or finalized at the Cancun meeting. Does that timetable work for everybody? Is anyone feeling that that's too aggressive or too lax?

Okay. So, let's sort of go with that. I think that's the last bit. Actually, I'm sorry, just as I'm looking in my email, Mukund, you sent a fairly long message to the list. Do you want to summarize that while we're all here on the call? And if you do want to summarize it, you would need to go off mute. Makund said he doesn't have a mic connected. Okay.

So, we do have one message on the caucus mailings about this. I will open a few more up to get this going. Like I say, I don't see any reason to let this linger long, but I also don't think we're in any rush. Fred, does that sort of timing work for you?

FRED BAKER:

That works for me. I was basically trying to figure out whether I should expect stable document and a vote by February or whether Cancun would be a better choice. And it sounds like you're telling me Cancun.

PAUL HOFFMAN:

I believe Cancun would give us the opportunity, because as some people comment on these new sections, it seems to take at least a week for other people to say, "Wait a second, I was thinking of this or that." So, that seems to be ... I think Cancun timeframe would be good.

FRED BAKER:

That works.

PAUL HOFFMAN: Great. And Ray just sent a message saying he's made some comments in the root zone hints file history. Great. Although that sounds like it's a content one as well, so we'll have to look at that in light of the whole thing.

I think we're having reasonable forward progress. I would love to see more activity on this. I would say the next immediate steps would be that staff will close out some of the comments that are in the document now that are trivial to do, just so that we don't have a huge stack of Google Docs, things on the right. And I'll get some things going on the mailing list.

Does anyone have ... Oh, Steve? Steve Sheng.

STEVE SHENG: Thanks, Paul. Can you hear me?

PAUL HOFFMAN: Yep.

STEVE SHENG: So, I want to ... Andrew raised this a little bit earlier. We were looking at the table of contents and determined how to call A-root. So, I think there's agreement there in both adding the organization as well as the fully qualified domain name of the root.

I just want to point out, in addition to just the top heading, there are other places, A-root, B-root, C-root, those are referenced throughout in the content of the document.

So, for example, if you look at section three, look at Verisign saying to operate A-root and J-root in a manner that [inaudible]. So, I think there's a lot of reference to things like a letter-dash-root. What do we do with that, given the table of contents discussion? Should we change all of that to the fully qualified domain name? I just want to bring this up front because this has editing implications. Thanks.

FRED BAKER:

Well, it seems to me that that can be addressed with appropriate global [replaces]. We can call it by the domain name or we can ... In this particular case, they're saying specifically that they would operate the A system and the J system as opposed to the things that Verisign operates.

It seems like at this point it would be good to use the domain name but let's limit ourselves to two names, the name of the company and the domain name used in the root service. Does that make sense?

STEVE SHENG:

Yeah, that's certainly [inaudible]. I just wanted to bring this to your attention. I think [inaudible]. I mean, in the text it says that A-root server and then in parentheses gave the fully qualified domain name. So, I think either that or in the early [inaudible], in nomenclature we

declare what nomenclature we use, so that we can just be consistent throughout the document. Thanks.

FRED BAKER:

Well, yeah, consistency is a good thing. I'm actually just thinking about this sentence, operate A-root and J-root in a manner that exceeds all RFCs and advice from related committees. I need to think about the word exceed. I think we're looking for something more like obeys or something like that.

But at any rate, the reference is that they're operating those [inaudible] Anycast systems as opposed to just the things that Verisign operates.

PAUL HOFFMAN:

I was not around when 023 was originally written. My understanding is that the subsections of section three are each pretty much controlled by the operator themselves. So, if an operator wants to keep the name, for simplicity, A-root, J-root, Q-root or whatever, that that would be fine. If they want to go with the full name, that's fine.

My understanding is that one of the reasons why the discussion in RSSAC of the going away from the letters discussion didn't come to fruition was that there was wide disagreement within the operators about what they wanted. So, I feel funny about us in this document forcing a change that may not be wanted.

Going back to Steve's question, I think that consistency is good, but not when you're forcing it on operators who like different names. And I think we can describe in words at the beginning of the document why

there are multiple names. We don't really need to worry about it until we get to section 2.7. That is everything before that, starting in 2.6 where we start using the names, doesn't have those names which is good.

So, starting with section 2.7 through the end of section 2, we can come up with a way of dealing with it, but we would be changing some wording that I believe was already heavily, heavily discussed during the preparation of the first version 023, especially in section 2.7.

We can reopen that. That's fine. My preference in order to get the document out is to, for section three, let the root server operators use their own names however they feel like it, as long as the reader can do reasonable mapping and that's our job to help them do that. And I think certainly a way to do that would be to expand ... If you look at the beginning of section three, just above section 3.1, Ozan, if you could scroll down to that ... Keep going, keep going, a bunch more. Keep going. Stop.

So, there's one very short paragraph there and I believe we can add to this. We can add a paragraph or two talking about the naming. And we don't have to name names—sorry, that's going too far—but to say some operators prefer one set of naming or another and to make it clear there, and then let the operators decide how they want their section written. Does that seem reasonable or does it seem more reasonable to try to be consistent so that the reader is always seeing a consistent thing? What do people think here?

STEVE SHENG: My sense is I think we have all the operators here and can agree on something that I would say this work party develop something and then have consistency because we aim for the reader in mind towards the end, and on that go, I think if possible we can reach a common naming nomenclature, that would be good. Thanks.

PAUL HOFFMAN: Wes, you had started to speak.

WES HARDAKER: Yeah. I was actually going to say the same thing that he just said, which is that I think consistency is better and we should be optimizing for the reader, not the writer, because there's a lot more readers than writers. I suspect—I mean, I can't speak for everybody but I think that coming to consensus is not going to be a difficult task and I think we even agreed earlier on naming that we'd want.

PAUL HOFFMAN: Okay. I wasn't there for that discussion. I had heard that there was not agreement. Fred? Saw your hand up, Fred. Sorry about that.

FRED BAKER: Sorry, I [inaudible] my mute button again. Assuming we have that consensus—and of course that's the first question, but assuming we have that consensus, I think there's room for an editorial hand going over this for consistency, kind of the way we have an RFC editor to make RFCs read [consistent].

What we could do, that we could designate somebody—I would suggest you—go through and make those changes, that Verisign [inaudible] J-root and so on and so forth. Then ask the operator if it's okay with them if that change is made.

I suspect that the operators will say, "Sue, fine, whatever." But they might not feel any sense of urgency in getting it done if it's just left to them to do.

PAUL HOFFMAN:

Okay. That sounds fine, although my preference would actually be for staff to do this since they've been doing a bang-up job already on going through and doing the universal search and replace. But we can talk about that.

And just to be clear, this is actually not a work party. RSSAC decided not to make this a work party. It's just sort of a document updating effort. I'm just talking because someone said, "Paul, why don't you volunteer to do that?"

Great. Any other thoughts here on this before we take a few things to the list, since we still have a bit of time on the call?

Okay. Well, thank you, all. Much appreciated. Ozan, do you want to close us out?

OZAN SAHIN:

Thank you, Paul. I'll go ahead and stop the recording. I'd like to remind all that we have another update session tomorrow at 14:00 UTC to update RSSAC 002. Thanks. Thank you for joining.

[END OF TRANSCRIPTION]