
MONTREAL – RSSAC Work Session 1
Sunday, November 3, 2019 – 09:00 to 10:15 EDT
ICANN66 | Montréal, Canada

FRED BAKER: 9:00 and all is well. Let's start out with a tour of the table, just people introduce themselves. Could we start with you, Duane?

DUANE WESSELS: This is Duane Wessels from VeriSign as the root zone maintainer liaison.

KEN RENARD: Ken Renard, Army Research Lab.

KARL REUSS: Karl Reuss, University of Maryland.

HIRO HOTTA: Hiro Hotta, WIDE and JPRS.

MICHAEL CASADEVALL: Michael Casadevall.

TOM MIGLIN: Tom Miglin, NASA [inaudible].

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.

WES HARDAKER: Wes Hardaker, USC ISI.

LARS-JOHAN LIMAN: Lars Liman from Netnod.

BRAD VERD: Brad Verd, VeriSign.

OZAN SAHIN: Ozan Sahin, support staff.

FRED BAKER: Fred Baker, ISC.

ANDREW MCCONACHIE: Andrew McConachie, ICANN org.

KEVIN RIGHT: Kevin Right, DISA.

NAELA SARRAS: Naela Sarras, IANA liaison.

JEFF OSBORNE: Jeff Osborne, ISC.

RUSS MUNDY: Russ Mundy, SSAC liaison to RSSAC.

DANIELLE RUTHERFORD: Danielle Rutherford, ICANN org. Support Staff.

FRED BAKER: Okay, great. So, one of the outcomes of the Metrics Work Party has been some issues that have been shown in these three documents; RSSAC2, RSSAC23 and RSSAC26. Andrew has been looking at those and has suggested some changes. So, this first session, this meeting, we want to take a look at 23 and 26, and the proposed changes and should there be any others, and just have a conversation about that. Andrew, let me turn that over to you.

ANDREW MCCONACHIE: Thanks, Fred. This session is about RSSAC23 and RSSAC26. The next session will be about RSSAC002. I put together a quick slide presentation, and I guess I'll just start off with that. Next slide? Thanks.

As Fred mentioned, the RSSAC is updating 002, 023, and 026. And then, beneath that, we've got what those new versions will be, so this will end up creating RSSAC002v4, RSSAC023V2, and RSSAC026v2. And for those of you who don't know all the RSSAC documents by number,

you might be like, “Well, what are these documents?” And those are the subsequent slides, but this slide is more of the overview of the work effort.

The RSSAC Admin Committee decided that instead of dealing with the heavy-handed approach of creating work parties and separate mailing lists and then having work party leaders and statement to work and all that stuff. This effort could be done a bit more easily, just a bit more less heavy and lighter if staff, I.E. me, wrote some kind of strawman text into the documents. And then, just facilitated discussions with the RSSAC pockets about these edits and eventually just worked these documents through the caucus until they eventually make their way to the RSSAC for a vote.

That’s the plan here with these three documents. There are guides, like designated guides from the RSSAC caucus for each of these documents. For RSSAC002, we have Ken Reynard who’s thankfully stepped up for that, and for RSSAC003 and 026 we have Paul Hoffman. And their role is to communicate as opposed to being leaders, like communicate updates and things of that nature.

I think it really is up to me to help facilitate the discussion and make edits and then take feedback from you guys and edit those edits. Now, of course, all caucus members are invited to participate, and all caucus members are invited to send text. Next slide? I think I have these ordered by number and not by session, so can we go down to RSSAC023.

BRAD VERD: I just want to make sure; I think you just said in your last slide, but this discussion is open to all caucus members?

ANDREW MCCONACHIE: Yes. Can we go to the slide on the section 023, I didn't order these correctly? No, you were in the right window, you just had to advance the slides a bit more. This is the first document we'll be talking about in this session. This update is, I mean, RSSAC0023V1 came out in 2015, I believe. And this effort primarily is focused on updating the history of the root server system, and all the things that have happened since then. There's been a couple big things like the IANA transition that need to get in there. Next slide please?

These are the proposed edits, and I'll go through these and then we'll take a look at the documents, and then people can comment on the specific edits. The first one is; each RSO has a section in RSSAC023, so we've sent out notifications or we've asked each individual RSO to update their individual section in RSSAC023, and we've gotten a couple of responses. The other thing we're going to be doing is adding text on the IANA transition, and that is an open action item to me and Staff, that text is not yet in there. And we'll also be adding text on Jon Postel's DNS root provision experiment. That text is already in there, and that's something that we'll be talking about today.

Another thing is describing the process for updating 023. Currently there's nothing in the document about updating the document in the

future. Now, the RSSAC does have guidance for updating RSSAC documents in general. Basically, just putting a V and then a version number after the V, and the RSSAC can update any of its documents if it so chooses. But there's nothing specific in RSSAC023V1 about how it should be updated in the future, so that's text we could add.

Finally, any kind of changes to spelling or grammar, we'll also make those. So, that's about it. Would it be possible for you to put the Google Doc in Zoom and we can just go through it? Brad, go ahead.

BRAD VERD:

I was curious, and don't want to jump the gun, but it seems like we spent four years working on 37, is putting 37 in there something that probably should be done, and maybe what 37 covers? We have that diagram that displays what the IANA transition was, and what 37 -- the scoping picture, something like that I think might be helpful. It's a big milestone, it was a document that a lot of people didn't think could happen, and it did.

ANDREW MCCONACHIE:

Maybe I jumped the gun a bit by diving into the document, and we should maybe have a discussion about other things that need to be updated first. That would probably make a good bit of sense, so let's start with that item, adding some kind of discussion about 37 and 38 into the history document. Are there any other opinions about Brad's suggestion?

That's a plus one from Liman, and a plus one from Jeff and from Kevin. And a plus one from Ken Renard. That's a lot of plus ones; I'd say, I guess my question would be, because that's still ongoing, where do we, what part of that history do we add? They're going to be kicking of the GWG pretty soon, so maybe, are you thinking just about the development of 37 and 38?

BRAD VERD:

Yeah, exactly, yes. Which I think you could -- I mean, there are two slides, so we gave the review to the GAC yesterday, I think there's two slides in there that I might grab and put in and maybe add some verbiage around, which one is the scoping slide that shows what 37, the governance that we're trying to cover. That blue box slide, and then the slide that -- maybe three slides. That slide, I think the principle slide, I think that is really important that shows the principles of the RSOs. And then, I would-

ANDREW MCCONACHIE:

The functions maybe?

BRAD VERD:

No, I wouldn't say the functions, it was -- let me look through the slide deck really quick, it'll come to me. There are three pictures that maybe we could add that just show -- oh, it was the workshop slide that shows the evolution of the document, what are the top picks that we went through, because that shows the amount of work that went

into the document from the RSOs. That’s just my thought. I’m not passionate about it, but I think it would be nice to put something in there.

ANDREW MCCONACHIE: I think that makes sense. And so, we could start off by including the slides and then include verbiage around the slides explaining them and then talk more about the history of the development of the document. But your idea is really to just focus on the development of the document; how they were produced, and the workshopping. Makes sense. Wes?

WES HARDAKER: I think that’s a great idea, and the way you can phrase it is; we saw a need back in, whatever starting year that it was where this whole thing started to come together, and it is still an ongoing process at this time. But the latest concrete steps are 22 and 37, and some of the things that came out of it. Don’t go into the GWD and other stuff, because it’s not that solidified, but you can say, “Here’s the last published elements of that effort.”

BRAD VERD: Just my train of thought; I think it’s an important piece of history around the root server system that shouldn’t be lost; it should be documented somewhere, and this seems like the right place to put it.

ANDREW MCCONACHIE: Okay. Go ahead, Russ.

RUSS MUNDY: Thanks, Andrew. I think an important aspect of this, and it's probably just a subtle sentence or phrase, is that the RSSAC noted that something like this would be needed in the future. Although other people had talked about it for a long time, that it was an RSSAC initiative that started this work that resulted in this.

MICHAEL CASADEVALL: One other thing I think needs to be added to this document is; we don't cover anything related to signing the roots, and that was a pretty big step forward to the history, as well as rolling the KSK and deploying IPD6 capable roots for 23.

DUANE WESSELS: I just brought up the document to refresh my memory on some of the things that it was about, and that topic is specifically excluded in the introduction, right? The report says it focuses only on organizational and social aspects of the root server system, and technical aspects are discussed elsewhere and out of scope for this report. I'm not opposed to including it, but that would be a big change. We would have to rethink the purpose of this document, I think.

MICHAEL CASADEVALL: Why did we do that back then? I don't remember.

DUANE WESSELS: I don't know. I'm just saying what it says now.

MICHAEL CASADEVALL: Either Andrew or Steve. Oh, Steve's not here. Steve was one of the primary authors.

ANDREW MCCONACHIE: Paul, did you want to respond to that? Is that why you're moved to the mic?

PAUL HOFFMAN: Paul Hoffman, for the record. I'm hesitant to talk about, to have a history document that's about what the root servers have been, talking about the content to root zone. And even though it was very exciting and rolling the root and the root server operators were acknowledging it, basically rolling the root was a small change, a small addition to the root zone.

So, I'm not sure that's appropriate. It's interesting, but it's not necessarily technically any more interesting than adding a couple of TLDs. And I don't think anyone would consider saying, "After 2012 the root zone got bigger." That doesn't seem that it would be part of this history document. And really, signing the root and rolling the root was the root zone getting bigger.

ANDREW MCCONACHIE: Did you want to respond to that, Michael? Go ahead.

MICHAEL CASADEVALL: Michael Casadevall, for the record. I understand where you're coming from, and I can definitely see the point. The main thing I would say, is that by introducing the KSK, we have another aspect of root zone management, because the key signing material has its own way it's managed and handled. So, I feel like not having anything about the KSK is important, because you cannot physically modify the root without it, so how the KSK was, the policy behind it came to be, and how the discussions led up to signing the root, if not the technical aspect in and of itself.

ANDREW MCCONACHIE: Wes?

WES HARDAKER: Thanks Andrew, this is Wes. I agree that there's elements of the content history that I think are worth capturing. My guess is that it probably ought to be a second document, because we probably ought to consider -- and Paul, I do think it's slightly more technically interesting than swapping a TLD or adding a TLD, but for the reasons that Michael pointed out. But then I think we ought to go back to what are the other important changes that occurred over time, such as, we set the TTLs to two days. Has it always been two days?

I actually don't know the answer to that, I'm sure it's in the back of this document, I think. But there's been a number of technical changes that have gone on over time, including the length of the descriptors as well as the meetings that were held to design, as well as the initial roll out of the fake key. I forget what that was called, but there's a number or interesting elements that I think are worth it. The DURRs, thank you, yes. That ought to be put in history before we forget it.

LARS-JOHAN LIMA: My devious mind just came up with the idea; why don't we ask the RSSAC to document the history of this content?

ANDREW MCCONACHIE: Is that a serious suggestion?

LARS-JOHAN LIMAN: Half and half.

ANDREW MCCONACHIE: Any more commentary on whether we should add verbiage on signing the zone or KSK roll? Brad?

BRAD VERD: Well, I just went back, I was reading the paragraph that you were looking at here, and we kind of contradict the first part of that sentence, if you look at the document. It says, "Notably absent are the

introduction of internet protocol version 6, addresses for root servers, the domain name, system security extensions and so forth and so forth. You can just do a search on V6, every letter talks about V6, so it seems like we contradict ourselves right there. But while I understand the content, certainly not adding TLDs, but it seems like signing the root was a major milestone.

I could go either way on the rolling of the KSK because I feel that, while I feel that rolling the KSK should be operational and should be no big deal, it wasn't. So, I think maybe the first roll should be documented and some of the things that went around it. And maybe I'm just saying that because there's no other place this is being documented; I'm not sure. But it feels like these are important pieces of history that should be captured somewhere.

PAUL HOFFMAN:

This is Paul Hoffman. I like Wes' idea of a separate document that is the history of the root zone, as compared to the history of the root server system. Because I think that once we open that up, I think we will find other things that the outside world had not heard of and considered, "No big deal," which were at least a big concern for people in root ops, if not this group.

And I think that that would be reasonably education to do that. And I don't think it would be a hard one to do, because you could actually almost do it sequentially. So, again, my concern is, and I agree the first paragraph started wrong, I just made a note that we need to

update that, but that separating the contents of the root zone from a document on the root server system, I think would be clearer to RSSAC000 that says, “We serve whatever is handed to us.”

BRAD VERD:

Okay, I follow that train of thought. I was just about to say before you said that, you could do another section within this document on the root zone, but then with that statement of the content of the root being somewhere else, maybe we shouldn't. Again, I could go either way.

RUSS MUNDY:

This is Russ. I would definitely fall in the approach that would say; address the content in a separate way so that, because especially in our current document, we have it where each of the operators goes back to their own beginning. But that doesn't, itself, particularly give an easy and accurate way to say, “What was the beginning of the root zone, and the root zone content?” And so, addressing them in a related but separate way makes a lot of sense, and in terms of responsibility, I would just off the top of my head suggest that it maybe should be a joint RSSAC and RZERC effort, since it is in both areas.

ANDREW MCCONACHIE:

It sounds like, what I take away from this discussion is that we need better text in the introduction of the document around scoping.

Regardless of what decision we make about whether this document talks about the zone information or the signing of the zone, the introduction is not very clear about whether this is a social history or a technical history. So, that could really be better clarified.

I'm also hearing that people, from Russ and Liman at least, that maybe a second document is needed on the history of the root zone, and that could be some kind of collaboration with the RZERC. That sounds like a big deal, and I'm looking at my chairs.

BRAD VERD:

Can I make a suggestion? Maybe we capture this list of things that we think we should capture, that maybe belong either in a different section or document, because we don't want to lose the ideas of things that we should capture, and at the very least, we can take the conversation up with RZERC.

And then at the very, very least, maybe we make an appendix saying that while we don't believe that this is part of the history of the root server system, we feel that these are specific milestones that should be captured, and in lieu of them not being captured somewhere else, we have them here. Something like that. But at the very, very least.

RUSS MUNDY:

Andrew, could you list them in a Google Doc so that people can go edit and add things as they need them?

ANDREW MCCONACHIE: List the different topics?

RUSS MUNDY: Yeah, the topics.

ANDREW MCCONACHIE: Yeah, not right now, but we'll definitely do that and send it out the list, yeah. That's an action item for Staff.

BRAD VERD: And if we're still talking about things to add to the document, again, this is for everybody, I know an e-mail was sent out about it, but the individual sections for your root server organizations, any updates that need to happen since the document was originally published, four or five years ago, should be done. It would be nice to capture any changes that you guys have done that are worth documenting in history.

ANDREW MCCONACHIE: I think Staff can send a reminder mail on that. Because we sent those out a couple weeks ago. We'll send reminder e-mails to all the RSOs on that. But if the RSOs could get their text in within the next couple of months, that would be really useful.

RUSS MUNDY:

It's worth noting, each of those sections, I just reread it, and 24 hours are written quite differently, and that's probably a fine thing. Each of the letters, some include history, mine doesn't for example. I'm now thinking, well, I probably ought to add that, because everybody else did. But it's probably worth rereading a couple of the other ones and see if you want to pull ideas to make it more cohesive. I don't think it needs to be stylized perfectly together, but I know I learned from reading from the rest of it.

ANDREW MCCONACHIE:

Good comment. Are there any other suggestions about things that need to be in this document before we dive into the document? Seeing none, let's go over some of the edits that have already been made to the document. We'll have the document in the Zoom room, but it might be easier if people loaded the document as well in their individual laptops. Starting off with the introduction, we've noted that we need, it sounds like we need better text around scope, and exactly what kind of history this is. I would like to skip down to section 2.8, if that's possible. That's the first new bit of text that we've added.

This is what we're calling DNS root provision experiment. This was Jon Postel's experiment in pointing things at him. This was text that was taken out of the original RSSAC23v1 and we've put it back in here. It's a bit old text. There's a lot of references in it. This is text that staff has written. I recommend people review it, and certainly make sure it doesn't contain any inaccuracies that you might be aware of and suggest comments. I don't know how many people have read it, or if

people are prepared to comment on it now. But if anyone has read it and has any general comments about it, now would be the time. Wes? I see you looking like you want to say something.

WES HARDAKER:

This is Wes. I did read it. I'm trying to figure out how to summarize what I'm thinking. I guess it feels kind of inconclusive in the end. Some things happened, there was actions taken. It was started and then switched. There are references, and you can go follow all the reference and try to learn more of the story. I think we had most of those references before. Even that doesn't piece together all the story fully well.

It feels like it just kind of ends, and there's a lot of whys that sit in your head, not understanding why things were done. The explanation talks about the whys for why it was done, but why it was ended or what his concluding thoughts were. It feels like it just suddenly -- oh, we're switching back without a whole lot of conclusion to an experiment. There's no summary.

ANDREW MCCONACHIE:

How could we provide a clean, I don't want to say happy ending, but how could we round this out in such a way that wouldn't just be interpreting history in a way that's possibly counterfactual?

WES HARDAKER: That’s why I was hesitant to say anything; I don’t know the answer to that question. I was trying to think of what I would do, and I don’t have a good answer, either.

ANDREW MCCONACHIE: Unfortunately, all we can really state is the facts. Go ahead, Duane.

DUANE WESSELS: I confess I haven’t read through this yet, but did you explain where this material came from? So, what’s the source of this; who wrote it?

ANDREW MCCONACHIE: I believe Staff wrote this. I would have to ask Steve that.

DUANE WESSELS: It’s my understanding that this was a controversial thing, and so maybe it’s a good idea to go back and interview some of the people who were interviewed for the first version of the history document and get their perspective on this. So, get a lot of perspectives on this event to make sure that it’s well represented, do you know what I mean?

ANDREW MCCONACHIE: Yeah, I understand. Go back and interview some of the people that were interviewed originally for this text.

DUANE WESSELS: Right and ask them if they remember this and if they were around at this time and what their thoughts on it were. I know some people who have opinions on this, and I can help you with that if you need.

ANDREW MCCONACHIE: Sure. I mean, I guess we would, get that list of people to interview and we can work offline, and staff conducted a lot of interviews for the first RSSAC023, and I'm assuming Steve did a lot of that work, and unfortunately, he's not here. I'm assuming that's where this text came from. We can definitely revisit that as well and go do that again.

I mean, I can work with you and other RSSAC members to get names of people to interview and maybe Steve still has his notes around as well. Yeah, that's a good action item for Staff to go conduct more interviews, and we don't need to sort out all the details of that right now, but that's definitely something we need to do. Paul?

PAUL HOFFMAN: Two comments on the text; one is, it's hard to read the text with a million footnotes. I'm not an academic, so I'm not a big footnote fan. Since almost all of them are to the same document, just different pages, just remove most of those footnotes and leave a footnote to that primary document. Because really, It's unlikely anyone who's reading this document is going to have a copy of the book open and then flip forward a few pages and such.

My other comment sort of going in a different direction from what Duane did but similar, and I wasn't at all involved in the first version, so I don't know it. I've heard opinions from people who were around at the time. A problem that we have always when doing history is talking about motivations. Saying what my motivation was twenty years ago is difficult. It's difficult enough now for me to say what your motivation is today, but to ascribe motivations for why certain events happened twenty years earlier is really, really difficult.

And I can, as a specific example in here, they talk about the IAHC, and as what they were doing, and that's the motivation for the IAHC. There were at least thirty motivations for the people in the room then and such. I think if we're going to do a history, we really, really should limit to; this happened, then this happened, then this happened. It'll make Wes unhappy, but the conclusion was; a thing happened. Not, why and anything like that. Or, I'm sorry, there are lots of whys but us saying it now I think would be significantly dangerous.

And again, this is from my very limited view, but I was actually part of the IAHC process, so didn't know about this at all; didn't care. But I can assure you, anything like that, and that would be the lighter part of this. Those of you who were here twenty years ago had a much heavier process, so I think it would be really important on writing this, even with the new interviews, Andrew, that you said might be being done, is that what goes in this document is very clean about events and not motivations. Thank you.

WES HARDAKER: Just to answer Paul, you're 100% right; nothing in there should be subjective. We should be restating what other things say or are known to say, not drawing new conclusions, absolutely. I wasn't trying to imply otherwise.

DUANE WESSELS: Yeah, I agree with Paul as well. My point was just that, since this was apparently a controversial thing, we shouldn't rely on just one source for its facts. We should get other peoples' recollections.

ANDREW MCCONACHIE: Michael, but first I want to ask Duane; do you think that in relying on these interviews we should then document precisely what the interview we provided with regards to the information in the document, so when we say something, say this happened and then this happened, we would then rely upon the interviewee as a footnote? Because throughout much of the document, the list of people that were interviewed are at the end, but there's no way to correlate the interview subjects, which with the information provided in the document. Do you think we need to do a better job of that?

DUANE WESSLES: I don't think so. I just think that some of the people that were around at the time and remember this incident aren't going to be paying attention to this document, and so I would just want to have them

read through it and see if what we're putting down matches their recollection of the way things played out.

ANDREW MCCONACHIE: And that makes sense, okay. Thanks. Michael?

MICHAEL CASADEVALL: Michael Casadevall for the record. As someone who came to this party a little bit late, I was ten when this happened, one thing that it is missing is context. For example, the document specifically notes that the generic top levels domains that existed at the time had a price increase but doesn't speak of anything of the other top-level domains, such as the country code TLDs which did exist. .US was added to the root in '85. I feel like there's a fair bit of context in which Jon Postel did this.

We've noted this in a few other places in the document, that X and X happened, but it doesn't say what the broader context of the DNS eco system was when that happened. Like, network solutions put these registration fees, but it doesn't say anything about other domains that were in there that would have influenced. We may not be able to draw a conclusion from that but being able to have that information in the same place might help put light. Allow people to at least understand what happened, even if we can't say explicitly because we don't know.

PAUL HOFFMAN:

I would strongly disagree with that in the sense that that is again talking about motivations and history that, to use Duane’s phrase, many people have opinions about. And as somebody who actually was active in the gTLD MOU at the time, it just seems for a history document -- maybe we can say something obvious in the document like, “Many people have opinions about this.”

I wasn’t in RSSAC when this section was taken out of 023 version 1. But the more context, as you call it, that we try to put in, I think that is going to turn into motivations for why did we get from step 3 to step 4? And such like that. Or, in your case, what you were asking is why did we even get to step 0? And I just think that that is easy to get wrong. In fact, it is impossible to get right, and I don’t think that that aids somebody who’s looking at the history of the root server system, knowing that bringing this up is a good idea. Because it was a part of the root server system, and it wasn’t just about the contents, it was about where the nexus of control was.

But going into all of the things around that will either misrepresent -- and I guess I’ll just say it; I’m very concerned about anyone saying anything about what Jon Postel believed. He’s dead, he cannot refute it. And that started happening about six weeks after he died. And twenty years later I think we should be super careful not to do that. But in the same light, a lot of the people who were active in this event are not around anymore, and we can try to interview them, but we won’t’ get it. I just don’t think that there’s value here.

BRAD VERD: Paul, just to make sure I'm interpreting you correctly; and I don't remember why this was removed from the original document. Are you suggesting we remove it from this document?

PAUL HOFFMAN: So, I would be okay with removing it. That is, when I saw it come back in, I didn't think, "This is super important." I think it has value, but it has value in a very limited range. So, if we leave it in, it should be as factual as it can be. And not as even chatty as the sections before, or after, about each root server operator. I wouldn't mind if it was a couple paragraphs that said, "This happened, then this happened, then this happened." And let the reader decide how important that is. But I also, if the group says, "Let's take it out," because it's too hard to write, I could totally understand that as well. Does that help?

BRAD VERD: I think so. I don't know, what's the view? Do people want to remove the section again? Is this not worth -- does it not provide value? That's really the question.

RUSS MUNDY: This is Russ. My recollection as to why we removed it before is that it was twofold; one, as it's been mentioned, it was a very controversial action that John took and how it was handled, and results and so forth. The action itself does not have a strong written historical perspective, because much of what occurred around this activity was

never written down in any formal way that you could get back to. And so, it was, in fact, extremely difficult to try to document the details of what occurred.

And perhaps Paul’s suggestion that a very short statement about the fact that this happened, it was an experiment, and it ended. Without going into any of the motivation or anything else, because I believe, and this is purely Russ Mundy’s opinion at this point, that it was primarily about who was going to control the content of the root zone. And that’s what is not available as far as a written documentation and we don’t, I think, want to get into raising that issue in a current today published RSSAC document. So, I vote, I would suggest that either an extremely short statement that it happened, and it ended, or that we just leave it out altogether.

BRAD VERD:

I think those are really good words. I would like to just point out that’s exactly why I think we should document some of these other major milestones that we don’t lose sight of them and end up in the same place. That’s just a side note.

WES HARDAKER:

If we are not able to document something accurately, then it shouldn’t go in, right? Maybe the right thing to do is say, “An experiment was done at this time, the facts are not well-documented, and we therefore are excluding it from this document. We acknowledge that something happened at the time, but the facts are not cohesive

enough to make a conclusion and document it as history.” I can’t talk today, I’m not sure why.

MICHAEL CASADEVALL:

Michael Casadevall for the record. I could definitely see cutting it down to what Wes just suggested, but I’m hesitant to remove it entirely, or at least, although we could move it to this new hypothetical document about the contents.

The reason I say this is that, I was actually unfamiliar with this incident up until I read this document, so I think there’s historical value in preserving it. We may have to cut it down to just the bare minimum facts and note that there isn’t that much information. I think not having it at all is rather important, because it puts information on how the root zone was, well, it was a major event in the root zone. That’s the one thing that you can take away from it, so moving it to a new document; sure. But I don’t think we should lose it entirely.

ANDREW MCCONACHIE:

Any other comments? Okay. So, what I’m hearing, please correct me if I’m wrong, is that this section needs to be cut down to be considerably shorter, to mention that the event occurred, lose a lot of the footnotes, maybe compress all those. And keep it very factual, mainly context free as we’ve called it, but in general this section needs to be considerably shorter, but not removed. Is that a fair reading of the room?

BRAD VERD: I think that's the next step, and then we can reevaluate with whatever the new, much shortened version is. And then, based upon what the feedback is there, you can make a decision on whether or not to keep it or remove it.

ANDREW MCCONACHIE: That's an action item to Staff then; shorten this section. I'm looking at the time, we're 45 minutes in, we have 30 minutes left. We're also supposed to talk about RSSAC026. I vote we continue talking about RSSAC023, I don't think we have very much left to cover, and then RSSAC02 probably shouldn't take, well, we're probably not going to finish everything in this session, but we'll just keep going.

I think that might actually be the only big chunk of text we have added new to the document besides lots of spelling corrections and reformatting issues. With regards to the individual RSO sections, I mean, there's still a lot of work that needs to be done in there; I don't think it makes sense to talk about that right now. There's an action under Staff to write a section on the IANA transition, that still needs to happen, so we can't talk about that. I believe, maybe I was wrong; maybe we can just jump right into RSSAC026 right now. Alright, let's do that.

RSSAC026, the RSSAC Lexicon. This is, again, an update effort. We're making version 2. This was published, RSSAC026v1 I believe was published in 2017 if memory holds. The RSSAC has started using

terms differently. There are also a few terms that we decided, or people suggested on the list that should be added to this. There's also a lot of work coming out of the RSS Metrics Work Party, which has changed some definitions in Lexicon. So, some of the definitions in 26V2 are dependent upon working finishing in the RSS Metrics Work Party, so we'll postpone talking about those, but we can talk about the other ones. Next slide?

The first point here, instance and root server that are two terms that have changed because of the work going on in the Metrics Work Party, and I think that work needs to finish and then whatever that work party decides, whatever comes out of that document, we're just going to mirror that in this document. Another change is, oh, Russ, go ahead?

RUSS MUNDY:

This is Russ. The concern I have with that approach is, we in the Metrics Work Party haven't really reached an agreement on how to say what we used to say when we said, "X letter," meaning the service provided for a particular letter. And we're sort of pointing at the definitions, the Lexicon change. That is something that is, in fact, mission in the current Lexicon document. I guess we could take that back to the Metrics Work Party, but I thought we'd sort of pushed it to Lexicon. We need to decide one way or another which way we're doing for defining a term that equals that.

ANDREW MCCONACHIE: I guess it really doesn't matter which work effort defines the term, because it's the same people in both work efforts, so yeah. To your point, yes, we need definitions for those words, and whether that happens in the metrics work party or in this work effort. Do people have strong opinions about that? Duane, kind of looking at you.

DUANE WESSELS: I don't know. What we talked about at our last Metrics call was to raise this issue in the caucus meeting this week. The thing, like Russ said, the thing we struggle with is, we really, when we're talking about things in the metrics document, we really are talking about letters, but we're tripping over ourselves to not use that phrase now.

So, I think one of the questions in my mind is, and I think we've been very successful in our efforts to distinguish letters from operators; I think that's been a very good change. But it's leading to this confusion in the Metrics Work Party because there we are specifically talking about queries sent to those addresses or those names. The question I would have is; are we comfortable now maybe, in certain cases, using the word, "Letter," again? Or, do we still want to totally avoid that?

ANDREW MCCONACHIE: Brad?

BRAD VERD: Can we not use identity? And it's the identity operated by such RSO? Or, is that not...

DUANE WESSELS: We can.

ANDREW MCCONACHIE: Duane, go ahead. That was a question for you.

DUANE WESSELS: I mean, we can. I feel like at some level, we just need to decide. And also, in some cases, it's weird things, it just kind of sounds funny to say certain phrases.

PAUL HOFFMAN: Does it only sound funny because of your historical reference to it all? Whereas if you're new, coming in from the outside, it wouldn't sound funny?

DUANE WESSELS: Maybe, yeah.

PAUL HOFFMAN: This is Paul. A problem that we discussed on the call with identity, is what we really mean, or even letter, is what we mean is; the two IP addresses that are affiliated with it, which I think would make something like that, which we already have a little bit of here under root server in the current one. But the problem we have for the

Metrics document is in fact, those two identities for the one letter are being treated differently, that we actually have different metrics for the V4 versus the V6.

It's going to take a good scrubbing in the Metrics document regardless of what we come up with to be clear that when we're talking about the operator that controls that identity, or however we do that, that's different than their V4 address. Or, we can't decouple the V4 and the V6 address in many places where we're talking about them, because we're explicitly trying to measure differences between the two of them. I believe it's tractable. I believe, Duane, you were saying it sounds funny, I'm pretty sure whatever we come out with is going to sound funny just from history, and I think it's going to be better.

RUSS MUNDY:

This is Russ. One of the challenges that we have is trying to look at this from the perspective, and someone who's not deeply knowledgeable with how the document is created and having them be able to extract a useful meaning from it. One of the points that was raised with respect to using identity as a generic term was that, as Paul said, if identity equals an IP address, each of the letters or the old letters, has V6 and a V4 identity. If you define, "Identity," to equate to the old letter designation, A, H, F, whatever, then we have one instance where an organization operates two identities, and it's very difficult to keep, have a consistent way for that to be read.

So, personally, my suggestion is that we at least in a few instances, maybe not as much as we used to, but I think Duane said this a minute ago, that we use letter again. Because this is, in fact, how the external world recognizes the usage of sending a DNS query off to move. It's listed that way in the DNS. A root server, blah, blah, blah.

BRAD VERD:

I'm not sure how to word this other than, I am confused. I feel like what you just said, the challenges you have with identity, I don't understand why you don't have the same challenges with letter. I don't understand the difference, I really don't. Especially since, when you look at, I don't know how many documents state the thirteen identities identified in the root zone. It's not thirteen letters identified in the root zone. I'm confused by the confusion. But I would like to understand that better.

The reason that the public, and again, I would say the public being mostly this community maybe, or IETF, identifies with letters is because historically that's what we've used, and we as a group agreed that maybe we wanted to change that. So, if we can avoid using letters, I think we should continue to do that. I'd like to understand the challenge, which I currently don't understand. I don't understand where using identity and letter are different. Or, using letter would solve something over identity, even when you talk about an organization that runs two. When we're talking about Metrics, we're not talking about specific metrics of a specific identity; we're talking about metrics in a generic sense, I would think. But I'm not sure.

UNKNOWN SPEAKER: If I recall correctly in the metrics, we are talking about the RSO and we have metrics per IP address, so no?

BRAD VERD: A letter has two IP addresses. An identity would have two IP addresses. So again, I don't see where they're different.

FRED BAKER: I personally object to treating the addresses as identities. They're how you access an identity, but in RSSAC2, we have separate statistics. This isn't new; the metrics work party. And they're accessing that identities' services, whatever they are. And by the way, they're UDP and TCP and other things that we play with, and we're reporting on those different access methodologies to that service. So, bringing up the two addresses I think just muddies the water.

RUSS MUNDY: Russ again. One of the things that has come up on some of the mail lists is that a number of people have said, "The way I interpret identity is that means an IP address, so every letter has two identities." And so, the term, "Identity," unless we, as RSSAC develop a consistent and hard usage of that term meaning exactly equal to what the old letter designation meant, in the broad sense, is needed before we can clarify. Because I was going to ask Brad, if that was in fact what you mean when you say, "Identity," that the identity that we used to know

as a root is the identity that you're talking about, as opposed to the A root IP addresses.

BRAD VERD: They are. They are referred to as identifiers, correct? So, I think identity is just the singular version of the identifier. So, there's 13 identifiers, and an identity would be one of the 13. That's how I think through it, maybe I'm wrong.

WES HARDAKER: So, what he just said, right? I don't have much to add to that. The one thing that's missing is, I'd love to see a sentence where identity doesn't work. And if I can see one, then I can better understand the problem. But the reality is that the word letter does not accurately describe what we're talking about. A letter has nothing to do with it; it is a named identity. So, the one word that we haven't thrown in there is the word, "Name." So, I'd suggest trying to do that, that we are using a name to define an entity which may have two IP addresses and I agree, absolutely, an identity is not a singular address. You're after three people, Brad. I think that's it.

LARS-JOHAN LIMAN: Lars Liman here. Two of you actually threw things around here. What if Netnod is serving only -- [inaudible] Tokyo, serves only IPv6, and our [inaudible] in Kuala Lumpur serves only IPv4? Does that change things? Probably doesn't. But I'm somewhat wary about introducing

terms that specify and nail things down to the nanometer position of exact things here because I'm not sure that's going to be helpful in the long run, when we talk to people outside this room.

I don't think I have a recipe for fixing that specific problem, but I'm not sure that we are helping ourselves by splitting the hair two more times. Identifier, identity, they're very close to each other. I'm struggling with finding a better word, but I really can't. Do we have to make this separation, is that important?

DUANE WESSELS:

The thing that we struggle with in the Metrics Work Party is, right now, the document uses RSO as an acronym. A lot of places when it really means identity. Or, it really means root server. So, for me, the question is, should we talk about root server identities or should we just talk about root servers?

And I'll give an example of where we find the language to be awkward. There's a section of the document where the way it says right now is, "Determining the number of RSOs required for reliable operation of the RSS." We could say, "Determining the number of RSIs, or identities, required for operation of the RSS." But to me, that doesn't capture the sense that those things have to be operational. Is it that they have to be operational or that they have to exist?

LARS-JOHAN LIMAN:

Can you read it one more time?

DUANE WESSELS: Okay, so the heading is, “Determining the number of RSOs required for reliable operation of the RSS.”

LARS-JOHAN LIMAN: Liman here. In that context, that’s actually two questions. One is; how many, using the old terminology, how many letters do we need? And given that number, how many organizations do we need to operate those letters?

DUANE WESSELS: In this context, it’s very much the first. And this is what we talked about back in our workshop in Rustin, so we spent a lot of time talking about this. But to me, it would feel more appropriate to use root servers in this kind of context. But if then we abbreviate it -- so, this is all editorial stuff, but if we abbreviate it, then it becomes, “The number of RS’s required for operation of the RSS,” which gets really icky.

BRAD VERD: Or the number of identifiers, which goes right along the line of identity. You could use identity and identifiers synonymously.

DUANE WESSELS: At some point we just have to decide and live with it, I think.

DANIEL MIGAULT: Speaking for myself, I was not confused at all. For me, RSO means an organization responsible to run a subset of root servers, and the question of identity never raised to my mind reading the document. So, only me. I'm not saying, I'm just providing my feedback.

DUANE WESSELS: Right, but it really is confusing. Even, for example, in the Metrics document at the end we have some example reporting and it says, "The RSO Metric for A root was 20 milliseconds." It doesn't say the -- and A root is not the RSO, VeriSign is the RSO. It doesn't say, "The RSO Metric for VeriSign," it says, "The RSO Metric for A root." I think if we continue with this current terminology, it's going to lead to long term confusion.

ANDREW MCCONACHIE: What I'm hearing is we have a couple of bad options to choose from, and there aren't any good ones. Do we just pick one?

DUANE WESSELS: Let me just bring it back to the terminology document. So, currently, we have a definition for root server, and we do not have any definitions for root server identity or identifiers as far as I can see. So, would we keep it as root server, or would we change that definition term to root server identity?

FRED BAKER: The identity and the root server are two very different things. My company operates 200 different servers; there's one identity. And confusing those two is going to lead to all kinds of confusion.

DUANE WESSELS: That's right, that's why we're here. And I would say your company operates 200 plus instances, not 200 root servers. That's how I would say it.

FRED BAKER: Okay, I'm willing to call them instances, or root servers; to me they're the same thing.

RUSS MUNDY: Our current Lexicon document defines a root server as an instance. And there is one of our problems in terms of consistency of use of meaning, because in the Metrics document, there are a few mentions of instance, but only to say, "The metrics document doesn't deal with instances, it deals with this thing known as A or F or whatever," that we don't have a good name for.

PAUL HOFFMAN: This is Paul. I have a question for Russ. Where in the current Lexicon document do you see something that says a root server is an instance? Because we should fix that; I don't see that.

RUSS MUNDY: It's the very first line of the definition, of root server. It says, "A root server is the name of an entry point, pare, an instance."

PAUL HOFFMAN: I'm looking at a live version of the document, I don't see what you're saying.

RUSS MUNDY: This is the published 026.

PAUL HOFFMAN: I'm looking at the V2, so we have revisions already that fix that.

RUSS MUNDY: If that's the set of revisions we want to go with. What I'm saying is, the current one out there.

PAUL HOFFMAN: I'm sorry, my apologies.

RUSS MUNDY: It says it's an instance.

BRAD VERD: We've gone way over on time. What we thought was going to be quick and easy, not quick and easy, and we need to spend more time on this. I think that's really the takeaway.

ANDREW MCCONACHIE: We should move on to a new term. This should probably not be so discussive.

BRAD VERD: This meeting ended like ten minutes ago.

ANDREW MCCONACHIE: No, this meeting goes on for another eight minutes.

BRAD VERD: Does it?

ANDREW MCCONACHIE: We have seven more minutes.

BRAD VERD: I apologize.

ANDREW MCCONACHIE: That's okay, we should move on. Because I don't think we're getting anything more out of this discussion. Does anyone have a problem

with lower casing the word ‘Anycast’? This should be a quick one. For some reason I remember there was strong discussions in the first version that anycast should always be uppercased. And I don’t feel anyone making that argument now. Okay, I didn’t expect that one to be controversial.

Moving onto 3; update root service to match RSSAC027. This is something that came out of a workshop, gosh, it was in like -- it was the workshop immediately after we published RSSAC026v1. We started using the term, “Root service,” differently. Let’s skip this one, because I think I can do the next three really fast. There are three new terms that are defined in RSSAC026v2. Does anyone have an issue with adding definitions for these terms? Not necessarily the specific definitions, but having these things defined in RSSAC026v2. DIDL, Root ops, and RSSAC. Those things that belong in the document. Liman?

LARS-JOHAN LIMAN:

I don’t mind, but in which context? What are root ops supposed to mean here? What’s the definition of it? Which root ops?

PAUL HOFFMAN:

Since I’m the one that added it. I actually lifted these directly out of RSSAC033. So, the document that said, “What is RSSAC,” and, “What is Root ops,” I hope I took the exact words from there. And I know that I took the name because I take the capitol O in ops for root ops, but I copied that faithfully. Let’s put it this way; if you are happy with the

definitions in RSSAC033, then those would be the ones that we would add.

ANDREW MCCONACHIE: I'll just read the definition for root ops that's currently in the draft document. "Root Ops is an informal community of all RSOs. It has no affiliation with other organizations, such as ICANN or the IETF. More information about Root Ops can be found in RSAC033."

LARS-JOHAN LIMAN: That works perfectly fine, thanks.

ANDREW MCCONACHIE: Before we close, are there any other terms that we're missing here? We have four minutes. Do you know of any?

WES HARDAKER: I don't have a term, but I did add a paragraph above all of the terms, trying to put a lot of the terms in context because it just felt like we were dumping into words without somebody having something to read, then to go, okay, I need to know what all of these things are. The goal was, they had this paragraph that they could read how it all flows together, then go read the terminology, and eventually go back and reread the paragraph to go, "Oh, it makes sense now." Delete it if you don't like it, it was a thought this morning.

ANDREW MCCONACHIE: Okay, you use the terms in sentences, providing context. Sure. Are there any terms that we're missing here? Okay. I think the really outstanding thing we have for this is to continue the discussion about instance and root server, if I understand correctly. And with that, we can break two minutes early. Oh, Duane, do you have something?

DUANE WESSELS: Well, just along the same lines, I was looking in my bookmarks, and it looked like at one point we started a document called, "RSSAC statement on identification of root server operators," but that was never published I guess, right?

BRAD VERD: That lost steam, and the person championing it didn't champion it anymore.

DUANE WESSELS: So, maybe it has steam now. I can send you the links.

ANDREW MCCONACHIE: I remember that document, but you're saying maybe we need more discussion. Not just discussion amongst RZERC and RSSAC caucus, but discussion like text discussion about instance, root server, identity, letter, all that stuff, in a larger document.

DUANE WESSELS: Yeah, as a standalone document, maybe that would be helpful to all that stuff.

ANDREW MCCONACHIE: Okay, so we I guess, Staff and the RSSAC Admin can look into possibly resurrecting that document. Thanks. Any other comments before we break? Seeing none, we have a break now, and we come back at 10:30, is that correct? Okay, come back at 10:30 for a discussion of RSSAC002v4.

[END OF TRANSCRIPTION]