PANAMA – RSSAC Work Session (4 of 5)
Thursday, June 28, 2018 – 10:30 to 12:00 EST
ICANN62 | Panama City, Panama

TRIPTI SINHA: Alright, let's get started. Welcome to Session 4 of 5, it's an RSSAC work session. This is an open work session. We have two agenda items for this meeting. One is an update on the RSSAC 000 document which is our internal operational procedures, and a KSK rollover plan.

I'm going to switch things around because we have a high priority agenda item. The ICANN board would like some advice from us on the upcoming ICANN updated KSK rollover plan. So with that said, I'd like to turn it over to Wes who's going to lead this discussion.

WES HARDAKER: Thank you, Tripti. This is Wes Hardaker from USC ISI. The ICANN board has asked RSSAC to review the current KSK rollover plan that will potentially be taking place later in October after it was moved from October of last year. As such, we have created a document to respond to this, and we're on a tight timeline to get this document finished by the end of July.

We are asking the RSSAC Caucus to help participate in this effort, so if you're an RSSAC Caucus member, you should actually have a link to a live, active version of our response in Google Docs and you can follow along. I'll go over sort of briefly kind of the conclusions we made at a very high level, but I'm hoping that most people have actually read the document who want to participate in the discussion. We do invite RSSAC Caucus members to listen to the audio stream and to participate in Adobe Connect, and if you have comments accordingly, please speak up.

Functionally, the document lays out in a structure that I'll reiterate as follows. Section 1 is the background and it basically states that on May 13th, the ICANN board has asked RSSAC to review the updated KSK rollover plan. Not just RSSAC, they also asked RZERC and SSAC as well.

In Section 2, we talk about the scope of our advice. We are technical engineers, many of us, and our desire is to analyze the situation completely and fully, but we are limiting the scope of our advice from RSSAC to be just from the viewpoint of the root server system. In other words, is the root server system going to be affected by the KSK rollover? And we are specifically excluding all other viewpoints. That includes what resolvers might experience or ISPs or end users. We believe that other

ICANN bodies with a wider remit will adequately address those issues.

Section 3, we dive into some known safe elements. One of the things that we wanted to do is reassure the ICANN board that there are certain areas that aren't of concern. Specifically, there are four areas, five areas in which we believe that adequate procedures and things are being taken into effect.

Really quickly, first off, we are reaffirming our commitment to serve the IANA root zone. We stated that in an RSSAC document before, but it felt prudent to do that again considering it will be IANA that will be doing this KSK roll at the highest level of the organizational tree.

Second, there will be no changes in packet sizes themselves, and therefore any changes to root server system traffic will not be reflected by packet sizes. We may see more requests – and we'll get to that in a minute – but the packet sizes themselves will not change.

Third, the RSOs have a commitment to OCTO to help provide real-time monitoring data of the root server system, and we've been doing that for nine months and absolutely plan on continuing that through this event.

Fourth, the RSOs are committed to doing DITL collection at the time of the rollover for post facto analysis and lessons learned type of approach so that we will have a historical archive of data that will be stored within DNS-OARC to make sure that we can analyze this event in the future.

And then fifth and final, we also pledge to be – the RSOs, IANA and OCTO will be working closely together during the event so that we can both monitor the event as it progresses as well as ensure that if a backout or a change in operation is needed, that we will be able to do so with efficiency and quickness.

One sec. So finally, Section 4 is a bigger section which is really factors that we believe that the ICANN board should consider in deciding whether to proceed with the KSK rollover plan as is, as I said earlier, which will be done October 11th if the plan is adopted – or I should say if the plan is not changed.

So first off, at a high level, the RSSAC believes that the reasons for rolling the root zone now is that the DNSSEC – what's it called? The DPS, what's that stand for?

UNIDENTIFIED MALE: DNSSEC Policy Statement.

WES HARDAKER: Thank you. The DNSSEC Policy Statement. I know acronyms, I'm horrible at expanding them. So the DNSSEC Policy Statement says ICANN will roll the key after five years. We don't believe that there's a technical need to roll this key from a cryptographic point of view. There may be other needs such as policy or perception type issues for rolling it, so that there's no technical need to press on as fast as possible.

One of the biggest concerns we have is the resolver software behavior is still somewhat unknown. Resolver software changes quickly over time, and there have been incidents in the past where a KSK roll by another section of a tree caused resolvers to send significantly increased levels of traffic. We believe a lot of modern software has been updated to deal with that. However, we don't know the deployment levels of that solution and we don't know… We believe that ICANN's OCTO department, the Office of the CTO, is adequately looking into performing some simulations, and we are suggesting that the ICANN board make sure that that is compete and the results of that simulation are taken into effect before considering the rollover.

We also want to make sure that all parties, including the RSOs, RZERC and IANA-related bodies have properly reviewed the published recoverability plan, specifically the KSK rollover backout plan which was updated in April of this year, and that

procedure needs to be in people's heads so that it can be executed as quickly as possible.

And then finally in Section 5, we have conclusions and advice which talk to the fact that we hope that these concerns will be adequately reviewed and addressed, but if they are, we believe that the root server system will remain stable and will not be affected by service outages and things like that if the KSK rollover plan is to continue.

We have a couple of outstanding questions by the people that have helped author this document to date. I'm going to take the second one first, which – the first one is talking about whether this is advice or not, and that's a good question, but I think we'll start with the second one first and then we'll conclude with talking about whether this is actual advice or not.

So Liman, would you like to state your comment or would you prefer for me to summarize it?

LARS-JOHAN LIMAN: Lars Liman here from Netnod. Please do the summarization and I'll catch on to that.

WES HARDAKER: Okay, fair enough. So I believe that, Liman, your statement which was attached to the sentence that reads, "RSSAC believes that estimating the increase of traffic load needs to be researched prior to the rollover, and the result of this research should be factored into the decision about whether or not to roll the KSK."

For the audience, please note that this is a working document. The wording is hardly final. We're trying to move this along as quickly as possible so the wording is not always perfect yet. This is very much an active working discussion.

So Liman's statement reads that, "I'm not sure whether we have consensus around this," and that we all need to discuss this as to whether or not we believe that traffic loads will need to be researched and whether they will be changing with respect to the KSK roll, and specifically with resolver software behavior."

I believe earlier in a working session with the original authors of this document, Liman, you stated that that comment was really reflected with the entire section even though it was marked with one sentence. Do I remember that correctly? Please.

LARS-JOHAN LIMAN: Yes. And as it stands here, it's wordsmithing where I would like to see the text not say "need." I would like to rephrase it to say,

"Further study would further enhance our fuzzy warm feelings about this." So if we could wordsmith that to something else, something less strict, I would prefer that.

WES HARDAKER: That makes sense, especially considering the conclusion text that you added earlier later which says something similar. Can I leave it to you to fix that particular sentence?

LARS-JOHAN LIMAN: Sure.

WES HARDAKER: Please do add the word "fuzzy," that would be wonderful.

LARS-JOHAN LIMAN: It goes slowly. When I need to find the right balanced English words, it takes me a few moments. So fuzzy was just what dropped into my mind.

BRAD VERD: May I suggest maybe comfort level?

LARS-JOHAN LIMAN: Very good. Thank you.

WES HARDAKER: Okay, so after you've done that, we'll resolve that one. Is there anything else that you want to address in a larger scope with your comment, which was really about whether we have consensus around the four points? And also, I ask the rest of the RSSAC members and/or caucus members to discuss whether or not we have consensus around the fact that there is no technical need to roll now, there are resolver software behavior studies that really should be done, and to review the published recoverability plans. So those are the three sort of points that we've put down to date. Are we missing anything or are there elements of that?

Joe would like to say something as a caucus member. I'd appreciate that. Thanks, Ken.

JOE ABLEY: Thanks, Wes. There are notes over here. I'll sit down. So just to put this in context, because we've already talked this week between RSSAC and SSAC, and I'm involved in SSAC in making a similar response. So I'm speaking now as a caucus member here, not as anything to do with SSAC.

WES HARDAKER: Correct, although I'd like to consider this an open discussion unless it goes off the rails.

JOE ABLEY: Okay, then I'll add a comment from SSAC at the end. So on this particular point here, if I was a nontechnical board member, I don't know that this gives me enough guidance as to be able to make a decision, and it suggests that this should be factored into the decision.

As a nontechnical board member, how do I understand what a negative impact of traffic would be? It says that the study in the increased traffic should be understood, but in the context of what? I think there ought to be more guidance here as to what might cause problems for the root server system and what would definitely not cause problems for the root server system, or whether you're talking about other relying parties like end users and you're imagining problems with middle boxes that might not take large responses. How is a board member who's not familiar with the minutia of DNS transport supposed to make a decision based on that recommendation? I think it could do with a little bit more fleshing out.

WES HARDAKER: I think that's a very valid point, so I'm adding a comment that we'll try and flesh that out later. That's a good point. Thank you.

JOE ABLEY: So I had another very minor point, is at various points in the document, you talk about at the point of the rollover. But it's not clear whether that means – which event it corresponds to. Is that the revocation of KSK 2010? Is it the point where the first zone without signatures made by KSK 2010 is published? Is it something else? It speaks as if there's like a definitive single point, and I think there are multiple points that could be interpreted as being the time of the rollover, and perhaps you need to be either more general or more specific.

WES HARDAKER: That's also a good point. We're referring to the timing of the rollover as not a continuous event of multiple subevents, but we are really talking to the October 11th date which is the expected date where the KSK used for signatures will change. And I think that we could absolutely clarify that. That's a very –

JOE ABLEY: Maybe just a footnote saying that's the date that you mean would be [inaudible]

WES HARDAKER: Yes, we could change the top text or the wording to be more specific to the fact that we're changing about the signature generation date switch.

JOE ABLEY: So then since no other hands are up so I feel justified in monopolizing the microphone, I'll switch hats then. One SSAC comment which is, just for the benefit of the other people who weren't in our direct meeting, that you have a series of reasons for rolling now, which I think are perfectly reasonable. You asked whether there's consensus on that. I'll switch hats back, that seems good. Switch back, SSAC hat.

I just thought I'd point out that one of the things that I believe will come out of the SSAC discussion is an observation that there is some operational value in rolling the key, and in particular, the observed problems with trust anchor distribution we think probably will go down in the future once there is an expectation that the trust anchor set is not static. As long as you don't roll the key, you're allowing that assumption that the trust anchor set is static to continue not to cause problems and become more ingrained. So SSAC, I believe, when they come to announce a consensus, that will be part of it. So just information form that side.

RUSS MUNDY: Yes. And as the other co-chair from SSAC, that work party, and involved in drafting here, I think this is the type of thing that at least in my model in my head for what we were talking about in this document, in the earlier parts of the document, it says other important aspects we address by other activities, I think that's kind of the category where I put that, because I don't see it as any kind of shortfall or conflict of any sort here.

BRAD VERD: Yes. I think just adding to that, I don't disagree with that sentiment, but I don't believe that that is in the scope of the root server system. That's why I believe it was left out here.

WES HARDAKER: Our original outline for all the things we wanted to include was significantly more extensive, and then we realized we needed to remit ourselves to just what RSSAC is principally responsible for. Thank you.

Any other comments from anybody? Especially concerning the consensus of Section 4 and whether we're all okay with it as is in terms of content and what we're saying.

Clearly, there are going to be changes made, especially based on Joe's comments and things like that. But it's a not document, we're not approving it today, but just making sure that we're all on the same page at this point.

Alright, so hearing nothing else, that really brings us to going back to talk about whether this document – so high on, we have a statement in Section 2 that says, "With that in mind, the RSSAC offers the following advice." And then we talk about safe elements and then we talk about our concerns, and the question that Liman had about this is, is this document advice, or should we change that word to findings or something else that doesn't say that it's advice?

That comment was made before we actually added sort of later in the document that actually, it sort of does create advice. So I'm going to read the three sentences in the conclusion from Liman's alternate text, not the text before it, that state that it becomes kind of more advice-related. And Liman, I did change your text a little bit to try and make it include advice. The very last paragraph before the outline. Yes, right there. So that text below the highlight.

"RSSAC believes that the current proposed timeline for rolling the KSK will not cause a service interruption to the root server system. RSSAC further advises the board to ensure that the

previous concerns are properly addressed before the KSK rollover. If addressed, these actions will strengthen our belief that the root server system will remain stable during and after the event." And again, Joe's statement about what event means is a valid one.

So the question is, I guess standing on the floor from Liman is, is our document considered advice, or do we only want to say that we're basically making findings? My personal take on this is that the board had asked for advice, so I think we would fall short it we did not provide it. Liman.

LARS-JOHAN LIMAN: Two things. I think you and I are slightly on different sides of the fence here, because I would like to recommend that we move forward. That said, what we're looking at are the technical details, and we have agreed earlier on that the reasons for moving forward are less on the technical side and more on policy side.

My reason for holding back is that I kind of want to avoid this advice to say, "Please don't resume the key rollover." And I would like to hear more opinions from the entire group of what you see, because I'm just one voice and I have heard few other voices with respect to what the hard advice that we're going to give with respect to resuming this or not.

The second thing is a nitpick that the current timeline is not going to cause any problem, but pursuing this according to the timeline will, so that's just a language nitpick.

WES HARDAKER: Yes, language things aside [somewhat]. I think you bring up a valid point, Liman, that we never really do state whether we should go forward or not. We state that we don't think that there are problems if certain things are taken into consideration.

One sec, Fred, I've got you.

So I added a sentence down below that we may or may not want to insert. And we're certainly not going to come to decision today even, we'll continue this on mailing lists and other discussions in the future. But at some point, do we want to include a statement that says RSSAC believes that we should move forward with the KSK rollover or should not move forward with the KSK rollover according to the timeline?

So Fred, Brad, and then Liman.

FRED BAKER: First off, I'm supporting Liman. The statement that RSSAC recommend delaying the KSK rollover from your text, that would say to me that you consider the policy side as kind of less

important. There's no technical reason, and I would agree with that, we don't know of a breach, but it kind of says that he policy issue is not important. And I think the policy issue actually is important. The reason to roll it, even if we don't know of a current breach, is that there might be a breach that we don't know about. And rolling it allows us then to obviate that. So it seems like the policy is actually important.

WES HARDAKER: Okay. Can I respond to that really quickly? Fred, if you want to change the text in that paragraph to try and make sure that we are not tying to state that technical is more important than policy, I think if that interpretation is what you got out of that paragraph, then I agree with you absolutely, that was not our goal. Our goal was to really document that there aren't bits on the wire reasons or cryptographic reasons for making a change at this point. So I'll let you look at that, but go on to your next point.

FRED BAKER: Well, I think my suggested text in doing that would be to go to Liman's text. I think they say the same thing with the exception of that factor.

WES HARDAKER: Which of Liman's texts? You mean in the conclusion?

FRED BAKER: No, Liman's alternate text.

WES HARDAKER: Yes, so I think we're already going to adopt that. There, we just did. Okay, I'm sorry. The alternate text – so Liman had a number of alternate text paragraphs in the editing session before we removed, and we failed to actually get to this one within the original authors. Based on the discussion we had earlier, I believe that Liman's text is what we will be using going forward. The only reason the other text was left was that the other text did sort of add advice, and we sort of agreed that we should be adding some of the advice. But let me go on to the other hands, which are Brad and then Liman.

BRAD VERD: So based on what I'm hearing here, I'm a little – how do I say this? I'm having a hard time going back and forth. And may I suggest, Wes, since you're moderating this, it seems like we don't have consensus on Section 4. Can we go back to Section 4 and run through each one at a time and see if there's consensus here? Because what I'm hearing over here is that reasons for rolling, like the DNS practice statement, the policy outweighs

the technical implications, so therefore we should do that. So I'd like to hear what everybody has to say around each of the bullets if that's possible.

WES HARDAKER: I think that's fair. I did ask if we had consensus for 4 and I heard silence and crickets, but I didn't – I think calling them out specifically is much better idea, thank you. But I'm going to go to Liman first.

LARS-JOHAN LIMAN: Yes, I fully support that. So let me see if my – yes, let's do that, and if my feelings are still not right, I'll talk after that. But I think what I was going to say would take us to exactly what you're about to do now, so that's fine.

BRAD VERD: If I can add one other comment – and Joe, maybe you can help me here – in the SSAC meeting, it was stated that they had no intention of stating do it or don't do it, they were going to point out things that should be done prior to it. And I think our job is to kind of point out risks, right? Point out the risks. Our job at RSSAC is to point out the risks and weigh in on if those risks should be investigated prior to it or not. And that it is up to the

board to make a decision as to whether to delay or not delay, whether or not they want to sign up for those risks.

LARS-JOHAN LIMAN: Can I fill in there? I would like to see us reach some kind of consensus and then make a statement that is – depending on where we end up on that discussion statement [that says] that we either think that there are technical risks that are so severe that we would really like this to not happen right now or that we don't see technical risks that are sufficiently high to block this at this point. That would be sufficient for me.

That's still not a recommendation to do it or don't do it, but it would kind of give them our assessment of the technical risks, whether they are so severe that we would recommend to not go forward.

BRAD VERD: But for instance, in the software behavior, we don't know how the software is going to respond. So are you suggesting that we need to get that work done before we give a recommendation? Because it seems like that work needs to be done so that we have an answer to then give advice. Does that make sense? Based upon how you just worded that.

LARS-JOHAN LIMAN: Yes, we could also see that his is a risk, this is an unknown. How dangerous is this unknown? Which are the endpoints? And can we say that even though we don't know, we still think it's okay to move forward because our risk assessment is that this will stay within a frame that is kind of acceptable for us.

BRAD VERD: I don't know if we'll reach a consensus if that's your point of view and mine is I think the risk is – I can't answer that question. You're saying you believe the risk is this big. I don't know how big the risk is because I don't have the data to actually prove that.

LARS-JOHAN LIMAN: Right. And I fully understand what you're saying and you are obviously right, but sometimes you have to move forward with something even though you don't have all the data.

WES HARDAKER: Alright, so let me jump back to a queue format because I think – no, I mean you guys were clarifying, so I wanted to get that out of the way, but I know that Joe wants to add something, and I do too at this point. So Joe.

JOE ABLEY: I think we are both addressing this from different perspectives, and I think that makes sense because I think RSSAC was really asked, "What is the impact on the root server system?" And SSAC effectively was asked a more general question. Even though the questions are formulated the same, the contexts are different.

So just because the SSAC is doing it some way certainly doesn't mean that the RSSAC has the same work to do. And to the degree that we have consensus today in the sense that we haven't heard any strong objections in the little work party that we have dealing with this response, SSAC has decided not to give advice about whether to proceed or not. it's in fact doing the opposite, it's being very clear that this is a decision that the board should make, and the SSAC does not want to give the appearance of making a decision about whether it should proceed or not.

We're giving references to studies, such as they are, like the study that OCTO's done on the 8145 data, the work to the extent that we can find things to link to that Wes has done, Jeff's various work and things like that, and we perhaps give some commentary about what we think our interpretation of those are in terms of risk profile, and we're also giving some guidance to say this is how we imagine would be a way that the board could assess the risk, understanding the risks of rolling and also risks of not rolling.

So we're trying to give a framework of how a decision might be made responsibly in our view, but we're going out of our way to say that we are not giving advice one way or the other as to whether it should proceed. And Russ, chime in if you think any of that's wrong or there's more detail that would be useful.

RUSS MUNDY: Yes. And in fact, the challenge is to scope the extent of the risk, because in a large proportion of the areas we're looking at in SSAC, the risk space is not well defined and certainly not precise. I think in the RSSAC context here, I think it's much more precisely to find both what's being dealt with and what's being pointed at. And in my view of the things that are currently in Section 4, the one that's the most challenging is the resolver software behavior.

WES HARDAKER: My bad. Alright. So from my point of view – so first off, Joe, if you need a link to my 8145 study and you can't find an adequate one, I will make sure one happens. I have a whole document that I've been meaning to publish as a technical report. I think I can trim it down because I'm not done with half of it. And I'll just do two, so if that's needed, I can get that out the door.

So especially to Brad I guess, if we changed our section that talked about the work that OCTO and the resolver study will come out with to say – because I think somebody brought up a valid point that we don't really say what to do with – I know it was Joe – that information once the study is done. It's like they could come up with some answer and it's like, "Oh, whatever the answer is, now we're happy."

And I don't think that's the case. I think if OCTO's team discovered that a load factor of greater than N% was going to happen, can we come up with a way to state I guess, A, what metric we hope to get out that study, and then B, a threshold where above this point we think that the rollover shouldn't happen.

And I don't know how to easily state that, and I don't think we can wordsmith that or decided upon that in this discussion forum, but would that make you happier, Brad, in terms of sort of stating the advice that we may or may not have to give with respect to whether we should go forward or not? Big IF. If we can do that, and I'm not sure we can.

BRAD VERD: Yes, sorry, I've got a whole bunch of things running through my head here. One is the whole advice piece. Are we going to give advice, are we going to point to risks and say you need to be

aware of these risks and it's your decision? Or are we making the decision saying, "Go forth?" It sounds like we don't agree on that, so I'll just point that out.

If we're talking specifically around the resolver software behavior, I'm not sure I can give you numbers. I'm not sure I can answer the whole question without having the data. I don't know what the risk is, what's going to happen if a number of small resolvers around the world stop resolving, what happens to the query load that send – some software in the past would send it to every root server, some would send it to a handful, some would – what happens? And is it a 2x increase, is it no x, is it 100x? I don't know. And 100x of nothing is not a lot, I get that, but what is the impact? I just don't know.

So I'm kind of like I have a hard time answering the question because I don't have the data. That's the issue. And I feel it's important data and it's not something – as we stated yesterday that you can't tell just by looking at the code. It was stated by OCTO yesterday that these are things that have to be done in a lab to see how they truly respond versus, "Well, the code says it's supposed to do this, so this is what we expect." We all know that that's not true.

WES HARDAKER: Right. Okay, Liman then me. And Russ.

LARS-JOHAN LIMAN: I can take a slightly different view on that. I'm with you, Brad, everything you said was kind of true. But you can also look at it from a different angle, which is that, okay, yes, we don't know what the effects will be so we don't know if that's 1x, 10x or 100x. But you can see that as a black box risk and say, "Okay, worst case it's 100x" or whatever the number is. It takes out the entire root system. That's a risk in itself. What are the remedies for that? How quickly can we back off? What are the effects of the root system going out for an hour or two until we roll back and get the proper zone – an old working zone place again?

So it becomes a piece of Lego in a large risk building thing. It's not the only thing that kind of makes the decision for us whether to go or not go. It would be different if we had data saying, "We know that his will happen" and it could be either good or bad. Then we would have more solid advice to give. But right now, we don't. so what we can say is either look into this more – and that's an open ended thing, you can look into these technical things to infinity and the cost will be equally infinite, or you can say at some point we limit ourselves and say this is an unknown risk, it's a black box, it's part of the entire risk package and that's for someone else to judge, but this is the black box that we know that we have here. And it's better to know that you have a black box than to not know it.

WES HARDAKER: Is that a response? Okay.

BRAD VERD: It's a response. So I feel like we're in agreement, but the terminology is off here. So let me be very clear, I am not saying work at this until you spend every dollar there is to identify what's going to happen in the wild. I think we all agree you cannot mimic what happens in the wild. But as you stated earlier, what he warm fuzzy is. It's getting to that comfort level, that, "Okay, we've done our testing, we've seen these things happen, we have an estimate – not a perfect number – of what the impact might be, an so our comfort level is this. We feel better."

Right now, I don't have a comfort level because I don't know what that is. So I'm not saying that we need to go out and answer some world where we test this and drive to make it look identical to the Internet. I don't think that's possible. But it's about finding a comfort level. Right now, we're saying – you just said that in your opinion, you're comfortable based upon what we have, which is speculation today. We don't know. And I'm saying I'm not.

So I agree with what you're saying in principle. You're saying your comfort level today is you're okay, and I'm saying I don't have enough data.

LARS-JOHAN LIMAN: Fair.

WES HARDAKER: Wait, I have a queue of people. Alright, so the queue just for reference is me, Russ and then Joe, and then Liman I think. Okay, and we'll go back to Liman. I'll delete your name from the top of the queue because you were in there, and then put you at the bottom.

Alright. So from my perspective, I think what I was suggesting before is, can we define the metric that we actually can measure against and give advice against? And a percentage increase of load from resolver X is not it. The metrics that I think we'd really want to see – and I'll get to it in a second in terms of calculating it with is what's the percentage increase to the RSS given a whole bunch of software out in the world that is increasing their load?

And really, that's made up of three different points of information. OCTO's team is hopefully going to estimate the load for particular resolver versions and hopefully other boxes,

and what they might exhibit under a failure case. I'm working on estimating what percentage of the load – and I started this yesterday but I haven't finished it – are we actually seeing from 2010 KSK data only? people that only believe in the KSK 2010 key. We know the load in terms of percentage of good versus bad keys, but we don't really know the load of how much of the server load they're actually accounting for.

Those are the easier ones to estimate. The harder one is the unknown percentage of – once OCTO's team comes up with unbound will increase by this amount, bind will increase by this amount, and maybe the CPE device will increase by this amount. We have no idea what the distribution is, and that's an unmeasurable number. So that's the biggest unknown in my factor, that, A, I don't think that we can even come up with a number for, and if we could, I think it would actually be a much easier calculation, but the reality is that we have no idea what the deployment is because we have no way of measuring what versions of things are actually being deployed.

So next was Russ, then Joe.

RUSS MUNDY: I think one of the possible ways forward we could go here with this is we've already cited in the document above the data collection for the loading and traffic and so forth on the RSS, or

at least 12 of the 13 designated systems. So that's being monitored very closely and has been for a while. It can be – and I imagine already is – being sort of looked at from what are the norms and what are the normal variations from these.

We've also cited in the document the preparedness for rollback if that's needed. And it seems to me that a way that we could move forward here with this would be addressing them sort of jointly. In other words, we know the traffic is going to be monitored, we know there'll be some change. We don't really know how much, but we should know at least a common amount of variation that's being seen today before the change, and predicting at what level we start to really worry if that traffic goes totally crazy. And then we look at, if needed, invoking the rollback plan.

So that might be a way to blend together the pieces we have to where we can say if these set of things occur, then there is a distinct set of actions that can be taken if required.

WES HARDAKER: Joe, please.

JOE ABLEY: Joe Abley, RSSAC Caucus this time. Just a very brief comment on Russ's thing. I'm a little bit concerned that what you just said,

Russ, is kind of conflating the results of the testing that OCTO is going to do on resolvers with observed behavior following the roll, because I think they're different scenarios.

I'll put that out there until – but the main suggestion I had was – as I understand it, if your timeline is the same as ours, and I think it is – well, the same as SSAC's which is not who I'm talking to, you're trying to give this response by August the 10th. OCTO was talking about having this testing done by mid-September, I think, in time for a board retreat.

I think the difficulty of it as we're seeing here is you're trying to speculate on the output of a study that hasn't yet been done. And I think as far as a workflow goes, what you might think about is to say we think it's important that the study is completed, we want to be able to see the results so that we can give a simple assessment to say, do the results of the study, we think, pose any additional problem for the root server system? And we'll give you a simple answer, either, "Yes, we think this study indicates there might be a problem," or, "No, this doesn't represent any kind of problem we can see."

That seems t o me to be a much more simple operational process, because they finish the study first, let's look at it, and then we'll give you a quick answer. Otherwise, you're in a sort of

Robert M. Pirsig situation of trying to define qualify of something that you can't measure because it hasn't been done.

RUSS MUNDY: Just for clarity, Joe, I was referring to the real-time monitoring of traffic, not the study.

WES HARDAKER: Liman and then me again. Anybody else?

LARS-JOHAN LIMAN: A different approach could be that – I kind of like Joe's things. I am looking at going back a bit on my previous statement of risks. So a statement in this document that said this is a risk that we right now don't know and that we cannot really asses. So this is a risk that we'll have to kind of bounce upwards at this moment because we don't have the data. That's a statement I could live with, and looking at Brad, is that something along those lines – would that be acceptable to you, saying that this is the most difficult item, we don't have any data?

WES HARDAKER: Go ahead, Brad. He asked you a direct question and I can wait.

BRAD VERD: Again, as I stated earlier, I feel our role is to point out the risks. Our role is not to make the decision. That's just my opinion. And I think it is – again, going back to, do we need to test every thing and have a perfect – perfect is the enemy of good, right? So I'm not suggesting that at all. I am suggesting that there is some reasonable level of data and information that needs to be done for each risk that might be identified. Does that make sense?

LARS-JOHAN LIMAN: It does. Yes. Let's move on. I think we have made our positions clear.

WES HARDAKER: I think you bring up an important point, Brad, which is that you do want to point out the risks and not make necessarily a decision ourselves. You want to give ICANN the ability to evaluate those risks though, and I think that that's where it gets harder, because we can point out, "Well, there's a sliding scale of percentage increase of traffic load that will greatly affect the root server system," and I think sort of starting from Joe's idea earlier, we are really considering two things.

There's simulation and estimation of traffic loads beforehand that need to be considered as well as actual loads being recorded during the event. And so the simulated, estimated one,

it would be nice if we could somehow give them a scale of risk where they're able to make that determination. If you just say, "You should consider the percentage," they have no idea what's a bad percentage of increase, right? So how do we as – good point.

How do we as RSSAC give them advice in terms of what percentage of load is bad? We're here, we each have appointed representatives from the root server organizations that hopefully at least have some technical clue in terms of what percentage of load is bad. The honest is we probably don't have an exact number to give them, but we could probably say that 1000% would be pretty bad but a 10% we could live with for a while. I don't know what those numbers are. But that same number would apply after the event. We need to provide them the ability to evaluate that risk. If we just identify risks and don't give them a formula for how to evaluate it, how are they going to make a decision?

BRAD VERD: Maybe I'm having a hard it me with a formula. And what I mean by that is I don't think – and I'd love to hear from – I mean there are more than three people at this table, and there are caucus members and whatnot. I'd love to hear from other people, but I

don't think we can answer the question with X percentage, I don't think we know that.

I think what you need to start thinking about is impact. Not this percentage of load increase is fine, it would be negative impact to the root server system. You're starting to see instances that are falling over, instances – some can handle, let's say, smaller amounts or smaller quantities than others while others can handle large amounts.

So the way the architecture is, it's not a single number, so I feel we should start looking at it as impact rather than percentage of load increase. Because certain instances might take very little increase in traffic and fall over. I don't think that's the case, but that's an extreme example that I'm trying to point out.

WES HARDAKER: No, that's valid. And just to respond, my percentage was an example. Remember I was trying to define a metric, that wasn't my proposal for it. Joe, then Liman.

JOE ABLEY: Okay. I was just going to say that I think that was as great example of why it's difficult to formulate in this document at the level that the board can digest exactly what the decision points are.

So just following on from the earlier thought that I hadn't developed that well, maybe just one of the conclusions of this thing is that on a specific point of the resolver testing, your advice on that is deferred until the results of that testing are available. And perhaps you could give a commitment to be able to give an opinion in 48 hours or something in order to provide some actual concrete timeline advice to the board and help OCTO make sure it's delivered in time for that, plus the board's advice.

Because it does seem sensible to me that OCTO making the decision about whether their study indicates that there's going to be damage to the root server system is kind of nonsense. This body exists to give exactly that kind of advice. It needs to be in that order.

WES HARDAKER: Carlos? Two questions for you. One, are you monitoring the room for any potential chatter in the messages?

CARLOS: Yes.

WES HARDAKER: Okay. Because I'm not. So I just realized, I didn't want to leave people out. And two, is there a board meeting that's happening late in September so that they could actually read and listen to advice?

CARLOS: Thanks. There's a board workshop mid-September, 14-16, around that timeframe.

WES HARDAKER: Does the workshop have a formal meeting associated with it that they can pass a resolution?

CARLOS: Typically, yes.

WES HARDAKER: Okay. Thank you. Liman.

LARS-JOHAN LIMAN: Yes. Thank you. Joe, I like your proposal to tack it on to the result of the study. Good thing.

Again, when it comes to risks and risk assessment, even if we feel that we cannot assess the risk because we don't have the data, it doesn't mean that other people can't. So by holding out

a problematic area to say, “This is something where we don’t have data so we cannot assess the risk of this” can still be useful to someone else saying, “Well, in the greater scheme, our assessment is that even if this blows up, there are other mechanisms that they’re willing to take a risk of the root server system going bad for half an hour, which could be the time it takes to back out to a functioning system again.”

And your impact thing which is spot on, absolutely, that turns it into a different type of investigation. Namely, how much damage can the root server system take before we start to see real effect on the end user side? And that’s a very large undertaking in my mind without having looked at it carefully. So while I think that is the right approach, I doubt that that’s something that we can do in a reasonable time. Definitely not before September or October, but I would think that that’s a multi-year undertaking.

WES HARDAKER: Alright. Russ, and then I’m going to propose hopefully a way forward.

RUSS MUNDY: I very strongly support what Brad was describing, an impact, that that’s really what matters. But the problem or the challenge

that I see with that is we don't – at least I have not heard of, and others may know about it, but I don't know how we can give any kind of quantifiable identification as to what that is. And how do you know when you've had too much?

And so this is one of the reasons for my suggestion of trying to use the data that we do have, which is the feed from the RSOs that come in centrally and all the RSOs can see what's happening with the other ones. At least from what little I've seen of it, that is something that is established, that does have all of the mechanisms set up.

Now, I don't know how much analytical processing has gone on or how much the results of what can be simulated in a lab can be added into sort of the expectations of what you maybe should be seeing when the rollover occurs and oh my goodness, it's just 1000 times worse or whatever.

I don't think we can get any more precise, because we simply don't have that much data. But if we use the combination of what we do know, which is the feed that we're getting, whatever analysis is done on it now, with what is expected variation caused by the lab testing, and then point to the fact if there are major variations from that, that's when the rollback effort or the rollback plan could be kicked into action if necessary. Because

the overall impact I'm afraid isn't quantifiable with what we have today.

WES HARDAKER: Okay. So I think that there is some consensus we can draw out of this. I think a couple of things. One, there's no way we can provide an accurate enough metric, and the best we can do is clarify stuff in terms of impact. I think that we all agree on that.

My takeaway analogy is that we'd like to state something like if the percentage of CO_2 in the mine rises above 8%, then – no, I said like – the miner shouldn't enter. And the reality is what we need to say is impact. So if the canary is dead, then don't go in. Right? We have to extrapolate it at some easier to identify level.

So we're not going to come to that conclusion today – and again, this document doesn't have to be done by the end of today, this is a working session to move forward. I think there's consensus that we cannot evaluate the risks related to the potential increase in resolver traffic until after the OCTO study is complete. And I didn't hear anybody objecting to that sort of language being inserted, and I heard more positive statements for it, so I think that we can go forward and change the text. Not live, but to state something like that.

So let me bring it back to the list in Section 4 that we talked about going back to decide whether or not – let me ask the chairs. Do we want to cut this off at some point so that you can do 000?

TRIPTI SINHA: Wes, no. Let's make this the focus. If you need more time, go ahead.

WES HARDAKER: Okay. So starting in the first Subsection of 4, which is the reasons for rolling now, I think there is consensus that there are no technical or cryptographical reasons for rolling now, but there may be political or policy reasons for doing so. We should not be weighting one of those two considerations above the other. Do we have consensus that – that summary, the words may not match what I just said perfectly, but that summary is accurate to everybody here?

UNIDENTIFIED MALE: I just want to know, do you want to remove these cryptographic things?

WES HARDAKER: No.

UNIDENTIFIED MALE: Because I'm not sure how related this is to RSS, because that's pretty much about the content of the zone more than...

WES HARDAKER: That's a valid point. Any response, Brad?

BRAD VERD: While the keys are related to the content of the zone, rolling the key and causing DNSSEC validation failures is directly related to the operation of the root server system. So we're talking two different things here so let's make sure we stay in focus on the root server system and impact to the root server system.

I believe my opinion regarding this – maybe this isn't worded the way we actually want it, but again, it comes down to risk.

So Fred, I'll look at you and I'll say I am not saying that the technical stuff outweighs the policy stuff, because that's not my job to make that decision. It's our job to point them out. We should point them out that there's no technical reason to roll. There is a policy reason to roll. However, the risk is that an exploit might come out tomorrow and you could be susceptible to it. But that risk exists today.

FRED BAKER: That risk could happen anytime, yes. Okay.

WES HARDAKER: So Brad, ready for the train?

BRAD VERD: I feel like I'm monopolizing the mic here for some reason.

WES HARDAKER: No, no, no. You have, I think, the best phrasing in your head for fixing some of this language, so I'm hoping to ask you to take a stab at addressing that in terms of risk.

BRAD VERD: Yes, I'm happy to do that. However, I think some of it might – this whole section might have to change a bit. For instance, in the first paragraph, the last sentence that you have highlighted essentially is related to the software resolver behavior. So I feel like we jump around a bit.

WES HARDAKER: Oh, yes. It was written an hour ago.

BRAD VERD: I understand, I've been there. I've been through this many times. I'm happy to go back and try to fit the words in there, but to me, it comes down to risk and us pointing out what those risks are.

WES HARDAKER: Yes, and I'll reiterate that I think today should be more of a high-level discussion, not wordsmithing. We will wordsmith over the course of the next month. And hopefully we'll set aside a block of time during the next RSSAC teleconference to discuss this as well.

Alright. So trying to move forward a little bit –

JOE ABLEY: Now we're on, good. So RSSAC hat. One thing that you go to some lengths to talk about is you're responding to a specific document, which is what the board asked. That document doesn't consider the question of whether to roll now. Not directly. It talks about rolling on a particular schedule.

So just on this paragraph and on the question of whether making commentary on whether to roll now is relevant to the question at hand, you could recast it to a discussion of whether you think the schedule that's proposed in the document that you are asked to respond to is appropriate in your view.

WES HARDAKER: That's a valid point. Can you put that in a comment on the document if you haven't done that yet? Brad.

BRAD VERD: This was a previous thought that I had when Russ, you were talking, I think – I don't want to conflate two other issues. One issue is response, which I feel we should be pointing out what the risks are. That's our job in an operational sense with anything you do with your company, point out risk. Right?

Russ, what you were describing in my opinion was an operational response to the role. So monitoring the real-time feed, seeing what happens and then making a call to roll back if it's needed. And that's an operational response. Where my concern is is saying, "Go forth, let's do this because we have this operational plan" without doing the reasonable due diligence ahead of time and trying to come up with some sort of estimation of what the impact is without doing – just like you do, you send software into QA, you test on it what's the impact going to be.

I think that's what I'm trying to suggest here without doing the due diligence in my eyes, taking an operational terminology, it's TIP. You're testing in production. If our only response is to watch

the real-time feeds and respond to them, we're testing in production to see if that's going to work.

RUSS MUNDY: I think what I was attempting to lay out was with the information that we currently have available to us and that we expect to have available to us with respect to the schedule, if it resumes on the time expected. And yes, it is an operational kind of approach, but the other concern that is really heavy in my head here – and this is sort of some of the bleed over from our SSAC KSK rollover discussions – is how much can we actually know?

So if we think we don't know enough to be able to make a reasonably sound assessment, then in fact we need to – whether it's laid out in terms of risk, identify it as a high risk that the board would look at and weigh in that manner. If we think that we want to give advice that causes the board to be very cautious, I think that is the way we should go.

What I've heard people discuss today and beforehand is that there are definitely concerns that there could be some big impact of one sort or another. But my general sense from the discussions is that probably the best thing to do is to resume the plan and the timeframe given the right cautions and so forth, but we need to be prepared to do the backout. That was kind of where I was trying to head.

BRAD VERD: An operational response, yes.

WES HARDAKER: When you pose risk, you give mitigations, and that's a mitigation.

RUSS MUNDY: Yes. Right.

WES HARDAKER: Okay. Nobody envies me right now.

UNIDENTIFIED FEMALE: It's a tough job [inaudible]

WES HARDAKER: It is. Thank you. Alright, I think we seem to be at a discussion ending point that we need to go revamp various sections of text. I think it's been a really good discussion, there have been lots of good ways forward and things like that, and there are good comments in the document now. So I think I'd like to declare success at this point in terms of discussion, with the caveat –

BRAD VERD: Success because you've thrown the pen over to me?

WES HARDAKER: That [would be] the caveat, besides the fact that I'm 15 minutes late to an appointment. No, it's okay. This is more important. With the caveat that, is there anything else that would benefit from face-to-face time? I think that we've talked about each of the sections fairly extensively to the point where I think we have a direction to go forward in.

The only remaining section in 4 that I haven't asked specifically about the small piece is review of the published recovery plans. I don't think that was contentious before, so I suspect we have consensus on that. But if anybody has a comment about the fact that we're stating everybody needs to be aware of that plan and believe that it's going to work – alright, hearing none. Is there any other face-to-face topic that we need to get out of this document? Tripti?

TRIPTI SINHA: This document? I was going to ask you if you were done rolling for today. Are you?

WES HARDAKER: That was just cruel.

TRIPTI SINHA: Rolling for today.

WES HARDAKER: She asked if we were done rolling it today. So it does feel like it's been rolled.

TRIPTI SINHA: Yes, exactly. Are we done?

WES HARDAKER: Yes I think we're done with this. Alright. Thank you, everybody, for your help.

TRIPTI SINHA: Alright. To be continued, right?

WES HARDAKER: Yes. And I apologize, I do have to leave.

TRIPTI SINHA: That's alright. Alright, the second session – why did Carlos walk out?

UNIDENTIFIED MALE: [inaudible] something else.

BRAD VERD: Do you have the data [inaudible]

TRIPTI SINHA: So I know Kevin's been keeping a log of what needs to be changed, and I see him online. Kevin, can you hear us? Are you willing to speak? Kevin? Alright, he's not responding. I don't –

UNIDENTIFIED MALE: He's typing.

TRIPTI SINHA: He is? Okay. Where is he?

UNIDENTIFIED MALE: It just says he's typing.

TRIPTI SINHA: Alright. Let's give him a second.

UNIDENTIFIED MALE: He says he's not on the bridge.

TRIPTI SINHA: Okay. He's dialing in. Let's give Kevin a couple of minutes to dial in. So this is about RSSAC 000, RSSAC's operational procedures.

BRAD VERD: I was under the impression that Carlos had a list that he had been kind of tallying. Yes. I just texted him to see if he can send us the list.

TRIPTI SINHA: Kevin, are you on?

KEVIN JONES: Yes, I'm on now.

TRIPTI SINHA: Alright. We've got about 20 minutes left in the session. Are you willing to run through your list of changes to RSSAC 000? And just to let you know, this is an open session.

KEVIN JONES: Yes. And Tripti, I actually don't have that list. I think – my understanding was for this session, we're going to try and see if there were things that we needed to add to a list.

TRIPTI SINHA: Okay. Alright. Kevin, you are echoing. Yes, there are things that need to be added to the list, Kevin, but I don't believe they were ready for this meeting, because we need to have that discussion at our next teleconference call. However, is there anything that anyone in this room would like to add to our operational procedures? Oh, you?

UNIDENTIFIED MALE: Makes sense Wes is gone. Wasn't there a point made yesterday from a caucus during the work party that our RSSAC documents say they are RSSAC documents and they don't say that they are RSSAC Caucus documents? Didn't we –

BRAD VERD: We did talk about it.

UNIDENTIFIED MALE: Talk about that and –

BRAD VERD: Resolved that.

UNIDENTIFIED MALE: That's resolved? Okay. Alright.

TRIPTI SINHA: Anything else? Anyone online got something?

BRAD VERD: Carlos has shared a few things with me. I don't know where he had to run to. I think he's probably supporting the other organization. He said documenting procedure for electronic votes. That was one thing that needed to be added. Updating the CSC terms to align with the CSC charter. That was the liaison position. And that's all I have at this second in time, but that was a couple of things that I if it spurs other ideas or things that need to be talked about.

TRIPTI SINHA: Alright. Anything else? Final call.

KEVIN JONES: If there is a decision to change what we're doing with our NomCom liaison in regards to voting, that would be a change as well.

TRIPTI SINHA: I believe that's on the list as well. I think that that actually goes back to us being appointed committee. I think that's a little bit more complicated.

BRAD VERD: Yes, I think that's a much bigger issue, and that's certainly not going to be addressed in this review of 000.

TRIPTI SINHA: Right. Anything else on your list, Kevin?

KEVIN JONES: No.

TRIPTI SINHA: Alright. Hearing nothing else, I'm going to declare this meeting adjourned. Thank you, everyone.

BRAD VERD: Yes, thank you.

[END OF TRANSCRIPTION]