
RICK WILHELM: Good morning, good afternoon, good evening, everyone. Today is the 17th of March, 2020. Welcome to today's regularly scheduled meeting of the RDAP RA and RAA Amendment Discussion Group. This is Rick Wilhelm from Verisign, hosting today basically in the absence of anyone else to host.

We've gotten regrets from a bunch of people, including Donna and Neuman and Beth and—who else did we get regrets from, Sue?

SUE SCHULER: Sorry. Jodi was a regret.

RICK WILHELM: Right, yeah. And maybe a couple others.

SUE SCHULER: Yeah. I'd have to go look back. Sorry.

RICK WILHELM: At least those come to mind. Graeme, Zoe says in the chat, will join shortly as soon as his other call breaks off.

We've got the document that we're going to be looking at. Sue pasted the link there in the chat. You should be able to scroll up in the chat, or you can just look right there. She repasted it again if you didn't see it. We're going to flip over to that document. Our goal here during this

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.

meeting is to go through this combined document and make sure everyone understands it and is in agreement with what's there because I believe that then our next steps that we'll be talking to—"It's important to begin with the end of mind," as various and sundry famous people have said ... So our goal here is to be taking this document ... I think we're going to be sending this out to our two stakeholder groups in order to get their general concurrence before we revert this back to ICANN Org as we present this to them. I think that that's our general direction. We can talk mor about that later, but I think that's our general trajectory here.

If I'm off in that, please stick your hand in the air and let's discuss that. But I think that's where we're generally headed.

I'm looking for hands. Don't see any, so I'm going to assume that that means we're largely in agreement with that.

Without any further introduction, why don't we go ahead and start doing this? Also, I think, at some point, probably towards the end of the call and certainly after Graeme gets here, I want to go through and make sure we put a [tickler] here to discuss a note from Rob Hall that posted to the list shortly before this meeting got going to make sure that we take care of discussing that important thing. That's in Any Other Business.

We do have Graeme that I see joined us here. Thanks, Graeme, for joining us. Graeme, there's the link that got posted in the chat. If you don't see that or don't have that to the document, holler and we can repost it to the chat. There it went. We're going to be going through this

with the goal of talking about how we circulate this to the Registry and Registrar Stakeholder Groups. Also, towards the end of the call, we're going to spend a little bit of time talking about Rob Hall's e-mail thread.

Anybody have any other things as we do a rough-and-ready agenda bash here? Anybody have any other comments on the agenda?

GRAEME BUNTON: No, Rick. Thanks for taking the lead and driving us forward. Apologies for being late.

RICK WILHELM: No worries. No worries at all. Congratulations on getting off that call.

Okay, here we go. Let's go ahead and tuck in. We're going to start on the RA section. The first thing that we can see on the document here—hopefully everybody can see that okay—is we got the stuff incorporating the changes to the registry operator monthly reporting. This says that, after this comes into effect, we're going to be adding something to the registry functions activity report. There's the additional of [cop] Field 38, which is the total number of RDAP queries responded. This has been what we've been talking about for a while here, so that should not really come as any surprise.

Anybody have any questions or comments about this one here on the reporting for the registry operators?

Very good. Seeing none, Sue is going to go ahead and give us a scroll. This is some stuff in the RA Spec 4 Section 1. What we're going to be

doing here is we proposed to add RDAP to the set of required services—again, not surprising—and we’re going to adding the gTLD RDAP profile as a set of requirements to do the actual implementation of that. Then this here contains the definition—the italicized language—of the RDAP profile. Let me get through this section and we’ll come back to Rubens’ comment in the chat. Here we’ve got language that specifically goes to the definition of the profile and refers to specific documents. Then it also provides the way for ICANN to describe specifically how these URLs are going to be updated. What this does is it allows the profile to be incorporated by reference, but it also provides a means for this thing to be updated via the policy update mechanisms. So it basically bring some certainty to the profile itself.

Can you scroll just a little bit farther, Sue? Okay. Now here with this one we also add a little bit about how the RDAP profile can be updated. We list specifically to update with the consensus policy and/or [temporary] policies and also through the cont[r]act amendment process. So this provides for no extra contractual mechanisms to change the RDAP profile. In other words, there’s no way to update it by “mutual agreement.”

So that’s where we are on this one. Again, I think that this is all pretty well-understood.

Anybody have any comments here on this language about the definition of the RDAP profile or about how it is being described to change it?

Very good. Seeing none, Sue, can you scroll back up just a little bit? Let’s go up to the RDAP reporting section. Rubens said in the that—Rubens,

you can come to the mic to elaborate if you want—we should define “responded.” So Rubens says, “We should define “responded.” Like, if HTTP queries do not match the RDAP specs or are invalid, should they include it or not?”

This one had more words in it previously. Actually, we chopped some words out of this in order to make it clear that we were only capturing the queries that were responded to. In other words, a query only was counted here if it was in such a shape that the service could respond to it. In other words, we don’t count malformed queries, truncated queries, queries that don’t meet some sort of a spec, or anything like that. They’re not counted. It’s only if the query is shaped and has content reasonable enough that it could be responded to. That’s why the wording is here. If the service can provide a response, then that’s what it can be counted as.

Jim has his hand up. Jim, breaking the monologue, please go ahead.

JIM GALVIN:

Thanks, Rick. I just wanted to clarify one data point to make sure we’re all in agreement here. The count of response is anything for which you created a response. It’s not about whether or not you actually provided what they asked for. If your response is actually a rejection, that counts.

Is that true? I just want to clarify that data point here so that we’re all on the same page.

RICK WILHELM:

I would agree with that statement.

JIM GALVIN:

I think [inaudible] should drop on the floor—I'm sorry, I didn't mean to jump in on you there—which I think are the things that you were talking about before. If you drop something on the floor for any kind of syntactic reason, or if that's your policy with respect to things you're going to reject, those don't count. It only counts if you actually send something back to the requester—whatever that is. Thanks.

RICK WILHELM:

Correct. That's why it's worded the way it is: the number that you respond to.

Rubens, does that answer your question or comment there?

"Yes, it does," Rubens says. Okay. Very good. Thank you, Rubens. Good discussion there.

Now we're moving down there. Go up just a little bit more. Right there. It says, "Define the RDAP amendment ramp-up period as the first 180 days after the effective date of the amendment." This is going to be useful later—that bullet there about the RDAP amendment—when we're going on and discussing the concept of the amendment period. So we'll use this a little bit later on. This refers to 180 days. We'll get down to this.

Now we've got a couple of things. This next one here on searchability is important here because this is something that we might get some pushback on: where the registries are stating that searchability does not extend RDAP services. What this says is that, if a registry has committed to WHOIS searchability based on a previous commitment, it's not

necessarily committing to that notion of doing RDAP searchability. In other words, when a registry signed up to do RDAP WHOIS searchability, it didn't make a forward-looking commitment to WHOIS searchability. We'll see how this goes over with staff.

Anybody have any questions about that one?

No questions? Catherine, please go ahead.

CATHERINE MERDINGER: Sorry. I was trying to find the unmute. My question was, will that mean they can't provide searchability? Or just that they don't have to?

RICK WILHELM: The intent is that they're not required to. It's—

CATHERINE MERDINGER: I'm just thinking, will we have to go through the RCEP process? Will the registry have ...

RICK WILHELM: Umm ...

CATHERINE MERDINGER: Because it's not something only a registry can provide, probably. Right? I do registrar stuff. I'm not an expert at [registries].

RICK WILHELM:

Right. The intent of it is, if someone wanted to add it ... Rubens says in the chat, "Searchability is provided only in web WHOIS." Right now in RDAP, there's not a searchability mechanism defined for one thing. The intent of it is that to be making it clear to ICANN that it's not requiring that a registry that is committed to searching something—searchability—in WHOIS is not automatically committing to something in RDAP. Therefore, I'm not sure what the requirement would be if whether ... I'm not commenting necessarily on whether or not RCEP would be required to add our searchability for RDAP because I'm not sure what that would look like.

Jim has his hand up. Jim, please go ahead.

JIM GALVIN:

Thanks, Rick. I think I just want to add a couple comments here in this discussion for us to keep in mind. I agree with everything you said. Quite succinctly I'll say that we don't want searchability to be required, but it should not be excluded. I like the specific point that Catherine made. We should make it clear that, as part of the discussion, we shouldn't fall into the trap of needing an RCEP added if we want to do it down the road. So I guess that that's one comment for us to keep in mind in our discussions—make sure that we all agree on that.

My second point here on this is, when I think about what ICANN is likely to say and what position they might have in this space, something to keep in mind is that, for new gTLDs in particular, some registry operators got their registries, their TLDs, in part because searching was a service that they offered because you got a point for that in the

application process. That was the way it worked under the old system. One wonders if there's any opportunity for ICANN to leverage that in a way to say that that's a problem from their point of view. If we're not going to carry that forward, then there's some kind of conflict there with the fact that they awarded it based on that and you're not taking it forward and it feels like a change of some sort.

Now, I don't have any answer to that, but I just feel like I want that data point on the table here in case others have a comment about that or at least we've heard it once. If ICANN raises it, we're just going to have to figure it out when we figure out what context [to] raise it in. Thanks.

RICK WILHELM:

Good point, Jim. Let's see. Rubens says, "As far as I know, no application can fail. If it hadn't scored the extra point for searching. Some TLDs have already moved searching from their agreements using the RCEP." Yeah. Also a good point.

Do we think—

JIM GALVIN:

But ...

RICK WILHELM:

Go ahead, Jim. Sorry.

JIM GALVIN:

Sorry. Just to respond to Rubens, he makes a good point, and I agree. But just adding to the conversation so that we've got all of our data here, it's interesting that he points out that they dropped searching via an RCEP. I just wanted to call that detail out because I guess, with what Catherine was saying, we need to be careful about how all of these pieces come together when we're actually talking to ICANN if we want to avoid an RCEP in this, which I think we should, because the technical side of this is going to be progressing. That's part of what's going on with searchability, and we don't want to be limited as we get to make forward progress. Thanks.

RICK WILHELM:

Okay. Very good. Let's see. I see Sam's hand up. Sam, please go ahead.

SAM DEMETRIOU:

Thanks, guys. I have probably a stupid question since I've mostly just been playing backup in this group. Does it really matter to us if adding searchability for individual registries is something that they would have to go through an RCEP to do? I'm just wondering if this is something this group really needs to be that concerned about. So, if there's a new version that doesn't have it baked into it, I would assume that means, like we covered before, it's still allowed. But it seems to me like it would be up to each individual registry operator to do their own assessment and work with ICANN to determine whether they would need to do an RCEP to add searchability back in. But I don't know that it's something that this group needs to spend a lot of time being concerned about. There's all kinds of things that registries could do that they would need

to use RCEPs to add. So I'm not really sure why this is a hang-up, I guess. But happy to be educated.

RICK WILHELM:

Good question, Sam. I can comment on that, but, Jim, you have your hand up. Please go ahead.

JIM GALVIN:

Sure. Very good question, Sam. I don't know that I have a solid answer, but I'll tell you there are two things on my mind from a practical point of view in this space. Certainly, if ICANN presses on this, we can take it back in our own group and decide our own consensus on the matter.

One thing that I have on my mind is just a practical concern. I want to be able to add searching when it's available and ready to be done. In today's world, there is no search spec that exists on the technology side. I don't want the RDAP Working Group to have to be limited in that space. As we move forward and as the IETF Working Group moves forward and a spec comes around, there are those among us that are going to add searching—registrars and registries alike. Why can't we just roll that out and make it a part of things? That would be a good thing. I'd like to be able to do that without any hindrance any anything getting in the way. So that's just a practical consideration.

The second thing that I would observe is that I think, in this space, ultimately, rather than doing an RCEP, that just makes it look a lot like a one-off thing and lot of one-off things. Maybe something for us to think about, if this does become a point of discussion, is, how do we turn this

into something that's more uniform? I would prefer that searchability, if it in any way is going to be included and put there, again, rather than an RCEP ... I'd like to find a way to make sure that we can do this in a uniform and homogenous way. I don't want to allow the option for a lot of variability in this space. That just doesn't feel like it serves the community very well either.

So those are just two thoughts. I don't have a definite answer, but hopefully that helps. Thanks.

SAM DEMETRIOU: Rick, if I could just respond quickly.

RICK WILHELM: Please go ahead.

SAM DEMETRIOU: That was very helpful, Jim. I'm glad I asked because I knew I was going to get a good answer from you because I always do.

Maybe what we should be thinking about then is something that looks a little bit more like opt-in to a single uniform service, as opposed to this language: "Update that section to state that it shall not extend to RDAP services." Maybe there's some room to be a little creative here where a registry could potentially opt into it without having to build the service from the ground up, like you said, once it's ready to be done, and not having to go through the RCEP process, which I agree is a giant pain in the butt.

JIM GALVIN:

I'm good with that. There's probably a way for us to come up with some nice words. Certainly we have some very fluent lawyers among us. There must be some way for us to suggest that, when the IETF has a specification, we agree that we're going to do it this way and that it's an open option.

The only question that then comes up is, is how the SLAs apply in that case or ... I think we know that the current SLAs don't apply to search queries. So then it'll become a discussion about whether there needs to be SLAs for searching and what to do there.

But I'm good with that, Sam. I like that. Thanks.

RICK WILHELM:

Very good. Thanks, Jim and Sam. I think that, here, one of the things that we're not ... Rubens made a point in the chat. He says, "If a future RDAP search is added to the fast-track RCEP, it's likely that the service would have low variability because most registries will prefer the fast-track."

I think this also is amplifying the rationale for why the current text about the current searchability not extending to RDAP is precisely because it's currently undefined and not standardized and not in a way that is consistent. So I think the points that the team is bringing up here really amplifies the validity of the point that we have right now that now is not the time for ICANN to be requiring searching/searchability to be moving from WHOIS to RDAP without some other

consideration/mitigating factors, primarily because RDAP searchability is so currently undefined/immature. So, while there might be a need for it in the future, whether it's a special-purpose searchability, as Maxim hypothesizes about it in chat: perhaps interest from local law enforcement or something, or whether [the] service that registries seek to provide that they want to be standardized as far as a fast-track, all of that is currently not defined. I think that all amplifies the point that now is not the time for ICANN to be requiring it. I think that we could get some general consensus on that.

So we might want to consider elaborating the wording around this point, but I think that the general point would be able to hold. We would be able to get consensus around that concept.

Any thoughts, either yeses or no's, around that?

Rob has joined. Rob has his hand in there air. Please, go ahead, Rob.

ROB HALL:

Thank you. Sorry I'm a bit late. I think, if we were to include searchability in any kind of discussion this way, the easiest way might be to just say the SLA applies per domain. If were to just simply put a word like that in now, even if we did get to a searchability discussion down the road, if we had to return ten domains, we get ten times the time. If we have to search 50 million domains, we get, in theory, 50 million times the time, which isn't reasonable. But it will focus ICANN on a "Hey, wait a minute. This SLA applies per domain." If we're going to talk about searchability or returning more than one, it's a whole other discussion.

RICK WILHELM: Okay. Very good. Thanks, Rob. I can comment on that, but I see Jim has his hand in the air. Jim, please go ahead.

JIM GALVIN: Let me respond to Rob and just say that my preference would be that there's no SLA I'm searching, that that's an entirely different service, an entirely different discussion. The current SLA, as it's defined, is defined explicitly for exact-match lookups. So there's an entirely separate discussion to be had about searchability because searchability is not just about how many domains you're looking up at a time. It's also about the size of the database. It's also about the number of fuzzy characters, if you will, that are in the query itself and all kinds of things. So my preference is no SLA because anybody who's doing searching is sitting there interacting and they should understand that. That's a separate thing, maybe, for us to talk about.

What I would prefer, in any case, is that we separate those two discussions. We have an SLA discussion about searching separate from what I think you're talking about here, Rick, in adding to this bullet item, which is we need to find a nice way to suggest that we're willing to leave the door open for searching to be included without actually saying exactly what that is.

I think the most important thing from a functional point of view—maybe I'll just type some words here in this doc for our consideration at the moment—is that it not be defined technically according to what's already in the agreements and then that we be willing to somehow

allow to be included when there is a technical specification. I don't know what the right way is to word that. I'll try to make those two bullets here so that people can think about it. We've got plenty of smart people here. Maybe somebody will have an idea of how better to word this before we share it. But that's what I wanted to say. Thanks.

RICK WILHELM:

Okay. Very good, Jim. Anybody else have any other comments that they care to make on this point?

While you're looking for the Raise Hand button, I would offer that, on the point of SLAs, I was going to say basically the same thing as Jim did, that searchability and SLAs are two things that don't really mix like peanut butter and chocolate. So we should keep those two things apart.

As far as something here related to searchability, we want to make sure, with whatever we put in here, we talk about something as it relates to optionality and avoid requiring on this because searchability right now, in terms of RDAP, is very uncharted territory. We really want to be careful regarding what we would putting into an amendment. So I think that that's really an important point here because not all registries are required to do this. The operational burden that we could be putting on ourselves could be considerable. So I think that's an important concept here. Especially in this post-GDPR world, searchability can have all sorts of unintended consequences.

Okay. Let's go ahead and try to keep moving because I want to be sure to get over to the RA piece also. Anybody have any other comments on searchability? For this one here?

Okay. Seeing none.

For the next one, Section 1.11, this is relatively minor. It's saying that the registries must link to ICANN materials and to lookup.ICANN.org. That seems relatively non-controversial.

Now we'll go over to the service-level agreements. Here in Section 1 [are] the definitions. We want to maintain the distinction between WHOIS and RDAP, each with its own definition, and scope those definitions in the context of Spec 10 Section 1 here and then make sure that the web-based lookups ... There is where we're having a web-based UI that only applies to WHOIS and not having a web-based UI that applies to RDAP. This is another position that we're expecting that ICANN is going to have issues with.

Here is the position that we've been holding: lookup.ICANN.org or similar clients is actually a unified web-based experience for people that are seeking access to web-based RDAP, which would be most of the users of RDAP. A unified experience is a better experience rather than having users be trained to go to the individual registry operator or registrar operator. We shouldn't be retraining users to operate in this decentralized mode of accessing RDDS information.

So what we're going to be talking about here with ICANN is, "Look, you've got this great tool: lookup.ICANN.org. As we go through this generational shift between WHOIS and RDAP, let's lean on that and transition to training users to go a centralized lookup portal and take advantage of the bootstrap mechanism that exists within RDAP and

train them to look at one spot for getting their web-based information.”
That’s really what this highlighted bullet means.

So that’s what this is about. Anybody have any questions or comments there?

Rubens replies, assuming he’s got a bit of a grin on his face, “A web WHOIS SLA could just be serving up a redirect to lookup.ICANN.org.”

Interesting point, Rubens. There’s a fair bit of truth to that. But that’s where this is. We’ll see where the discussion goes with ICANN on that one. This is really, I think, the place where we should anchor our discussion on: this is a better user experience rather than having users go to individual registry operators for all the data.

Moving on to the SLAs—now going on to the table here—we’ve got availability at 98%. This is the WHOIS queries at 2,000 milliseconds, which is the same WHOIS update at 60 minutes. Then we’ve got the RDAP queries. You can scroll just a little bit, Sue, to where we see the RDAP information. RDAP availability is at 98% also. Then this is the one that we know that ICANN is not happy with. They don’t like this 5,000 milliseconds. Russ has intimated that he wants to get this back down to 2,000 milliseconds. Both Jim and I and others have had a number of discussions about why there’s not a real basis on that for 2,000 milliseconds. So we’ll see [where] that turns up.

I see Rob’s hand is up. Rob, please go ahead.

ROB HALL: I think we have to push Russ on this a little bit. He's getting the 2,000 milliseconds from the registry contract. I know it's 4,000 in the registrar. So I think they're just stuck on, "Well, wait. It's 2,000 now. We want 2,000." I agree it's much more complicated. We need to educate them on that. But I think we need to push them off that. Right now, for an end user, it's actually 4,000 milliseconds, not 2,000, because most go to registrars to do it.

RICK WILHELM: That's a good point, Rob. Thanks. I agree. In these days, the 2,000 milliseconds really doesn't serve any people that are mining data. So that's another concept that we'll lean on. So we're going to be leaning and pushing on 5,000 milliseconds. We'll stick there for right now.

Anybody have any questions about—

ROB HALL: [inaudible]

RICK WILHELM: Sure. Go ahead, Rob.

ROB HALL: One more. I would also point out that I don't know of any registry that really sets their clock now to, "Oh, I have to do it in 2,000. I'll do it in 1,950." The way these systems work is that they respond as fast as they

can, and only when under extreme load do we ever get anywhere near that.

RICK WILHELM:

Well put. JC says, "Couldn't we offer 4,000 milliseconds as a compromise? It meets the actual situation and is lower than 5K."

Good question, JC. What we haven't yet seen is anything from ICANN that indicates that they're willing to stop at 4,000. We had talked about and settled on 5,000, and they had agreed to 5,000 and said 5,000 was fine well over a year ago. Now they've said 2,000. So we're in principle sticking to 5,000 because they had agreed to it and they had no reason to go back and start talking about 2,000.

I see Brian's hand is up. Brian, please go ahead.

BRIAN KING:

Thanks, Rick. Just an advocacy approach, I agree with JC. I think JC's approach does help keep them honest. You said there's no indication that they're going to stop at ... I think it was 2,000, you said. If we had to position this as a transition SLA, we could have a tool in our toolbelt to say that, for the first six months or the first year post-implementation or post-agreement, it could be 5,000, but, after that, it would be 4,000, but no lower—something like that—as a sign of good faith. But let's just keep that in our back pocket. I think you're right, Rick, that we need to put the onus on ICANN to take a step in our direction. Thanks.

RICK WILHELM: Very good. Thank you. Good points, Brian. I see Jim's hand is up. Jim, please go ahead.

JIM GALVIN: Thanks, Rick. As Brian said—"keep in our back pocket"—I think the key thing to pull out of discussions that we've already had with Russ on this point is that, before we consider anything less than 5,000, we want some evidence, and we want to collect some data. As operators, we just really have no sense of what the performance is likely to be on this. So 5K just feels like a good place to start. It ought to be far enough out there that nobody is going to run any risks. At some point in the future, when we get some more evidence that it's all good ... Because, as you said—well, I don't know if it was Brian or Rob; I forget now; I apologize—we're all going to provide the service just as quick as we can anyway and probably meet the SLA. So it's not really about that. But we want to be careful to have some actual evidence and actual knowledge and really data-driven solutions here. That's the real key thing to be pressing on in this. Thanks.

RICK WILHELM: Very good, Jim. Thanks. Good point. The other thing to that end is that volumes are only going to go up because obviously most of the folks within earshot here are running RDAP right now, but volumes are only going to go up and usage is only going to go up as adoption grows. So the numbers that we're all getting now are only going to go up. Of course, in order to meet and improve upon SLAs, that costs money. This is not exactly the revenue-generating side of the business for anyone

here, as you all know, as he holds his hands up in the gesture of surrender. So this group clearly knows that.

Jim, please go ahead.

JIM GALVIN:

Thanks, Rick. Sorry. I'm just chuckling a little bit. The numbers are only going to up, but the other thing to keep in mind is it's not just about RDAP. We're all fresh off of listening to EPDP updates and discussions and GDPR stuff.

One of the things that will have to be ultimately be clarified here is, where is the validation of the query? RDAP is fundamentally different from WHOIS. You've got authentication to deal with now, and that kind of volume is going to go up, too. The operational cost of that is probably included in this query SLA because presumably it starts when the query is made on the client side, and it ends when the client gets a response. That means that all that is shoved in here, too. We really have no clue about what SSAD is really going to turn into or whatever it evolves into as part of this process.

Sorry. Too many unknowns here. We need some more data. We need some more evidence. Then we're willing to talk about what might be better if it really needs to be. Thanks.

RICK WILHELM:

Yeah. Good point. Now, I think, though, that all the SSAD stuff is outside of this bucket. Is that correct?

JIM GALVIN: Well, again, this is a discussion to be had. You're right: You have a hybrid approach where all the authentication and validation stuff is done at ICANN, and then the query with whatever hints or influence are provided to the registries and registrars for the response. Yeah, all of that is outside, but if any of the authentication has to come in to the registry or registrar—any of those tokens ... So it just depends on where the pieces fall, right? Thanks.

RICK WILHELM: Okay, yeah. Because I thought that all we were talking about here is that all this stuff that we see on-screen was only for publicly available RDAP.

JIM GALVIN: Yeah, you're right, Rick. You're right. We don't actually use those words anywhere up here, do we? That's probably a point worth making very clear somewhere. That's important.

RICK WILHELM: Yeah.

JIM GALVIN: I'm going to make a comment up here at the top. I'll put a comment in to note that.

RICK WILHELM:

Yeah. "Needed [as for] public," as Maxim says. Okay, good. That's good for us out there because all this stuff is similar in the way that searching is outside of the SLAs. Anything related to SSAD is also outside of the SLAs because that's a different kettle of fish, as the saying goes.

Let's keep going here on this. I think this is good discussion. Now we've scrolled down a little bit. Then we've got some other comments here at the bottom. "Registry operator be required to meet the RDAP SLRs defined above." "Update Section 4: The [changing of terminology of RDAP to WHOIS services as applicable." Then, "Add a new] Section 6 to add the specific of the SLR." Then, "Changes [inaudible] Section 7 on emergency thresholds to update and define the thresholds during and after the wrap-up period."

What this section here basically means is that this has us avoiding double jeopardy. That's what it amounts to: when RDAP and WHOIS are both operational, you have an emergency threshold trigger only when both RDAP and WHOIS are down during the time when they're both operational. It says, "Any failure to meet the service levels set forth in the service-levels document shall not trigger the WHOIS emergency threshold unless both RDAP services and WHOIS services have concurrent downtime and each has exceeded the emergency threshold." So this says that, basically, registry operators are taking on an additional RDDS, and the RDDS emergency threshold trigger does not get triggered unless both your WHOIS and your RDAP are both down simultaneously and each has exceeded the emergency threshold. So this avoids double jeopardy on the emergency threshold triggers.

Does everybody parse that?

I'll take silence as a yes. "Rick, your explanation was brilliant"—that's what silence means. I can see Graeme grinning, even though Zoom doesn't have video working. Very good.

So that's the end of the registry section. Anyone have any comments on the registry section before we go to the registrar section?

Nothing there. Okay. Let's move on the registrar section. Let's start here. The first one says, "Registrars are providing RDAP for all gTLD registrations sponsored by the registrar." This says you provide web-based WHOIS as you currently do, and then data from RDAP must be from the registrar's own database.

Let me stop right there and make sure that we're parsing all of this because this is a macro thing on top of this because I think that this gets to the thread that Rob had talked about a little bit earlier, just before the meeting was getting started. Registrars are doing RDAP for all gTLD registrations sponsored by the registrar, and it's from the registrar's own database. That's the opening salvo here on this one.

Rob, please go ahead.

Still on mute, Rob.

ROB HALL:

My apologies. You're right. I'd love to be able to add something here: in an emergency, I can fall back to the registry's database. It seems pretty exact for no reason if they're going to hold us to this. If I'm having an RDAP failure, why can I not proxy?

RICK WILHELM: If you're having an RDAP, why can't you proxy?

ROB HALL: Or, if I'm having a database issue—I'm taking my database offline for upgrades—why am I not allowed to use the registry's source? If it was language such as preferred ... I get what they're trying to get at here, but offering a result is better than not. If I can't provide the data, why would I not use then other source? Or why would I be prevented? It seems overly constrictive.

RICK WILHELM: Okay. Anybody want to comment on that one?

Rubens says, "Because the registry might not have the data. This might work or might not work on a per-TLD basis."

ROB HALL: Rubens, I'm just talking about ... This seems very narrow and restrictive and restrictive and ties our hands as to how we provide the service. I'd like to get rid of it altogether, but, if they're not going to go for that, why am I not allowed to use that source of data as a backup?

RICK WILHELM: Jim, please go ahead.

JIM GALVIN:

I'm just thinking through that a little bit, Rob. I think one of the observations I would make, just thinking out loud here on all of this, is, in addition to what Rubens is saying, there is clearly the question of whether you would be subjected to any kind of limitations, even at the registry. In fairness, is a registry likely to rate-limit you in any way or find that abusive? Or would you need some kind of special arrangement to do that? I'm just thinking out loud on options here and things that might happen in addition to what Rubens is saying.

It's also possible that the registry itself—yeah, as Rubens is saying—might not have the same data or might be under a different jurisdiction and thus give a reply that's not quite the right one. So it's probably more accurate for the client to know that they're not getting the data for you as opposed from the registry.

Brainstorming a little bit more—I'm sorry; one step further as I think through this—if you're going to do something like that in an RDAP sense, what I think you ought to do as your fallback is send back a referral to the client to tell them to query the registry so the client knows what it's doing. So you shouldn't do the recursion of what you're doing now, but take advantage of the RDAP protocol and the way that it works and refer them to the registry if that's what you want. So your fallback position ought to be that, from a continuity point of view, you have the ability to stand up a server that just does referrals back out. That, I think, is the right answer. Thanks.

ROB HALL:

Can I ask one more dumb question then under this section.

RICK WILHELM: Sure. You can ask any kind of question.

ROB HALL: I notice that the first two bullet points talk about RDAP and then the [third] one talks about web-based WHOIS. What about web-based access to the RDAP?

RICK WILHELM: Good question. Is that captured somewhere else? Here I'm scrolling. You don't have to, Sue. I think that's an omission.

ROB HALL: Because I'd be happy with it omitted if it means we have no SLA on web page.

RICK WILHELM: Wait a second. Is there a requirement for web-based RDAP or web-based WHOIS for registrars right now?

ROB HALL: Yes.

UNIDENTIFIED MALE: Yes.

RICK WILHELM: There is? Okay. So then I think that, for consistency between the registry and the registrar, there would be a concept that we would say that we would grab the paragraph that we have up above in the registry section—oh, it looked like ... I think I saw Jim just copy and past that into the comment. Yeah, that's what he's doing. He's doing it in real time.

ROB HALL: [inaudible]

RICK WILHELM: Yeah. So it's in there now. Jim just added it.

ROB HALL: I want to make sure I can understand what they're asking. So a registrar no longer has any requirement to have a web page with any kind of customer data on it? I don't think that's what they're thinking here.

OWEN SMIGELSKI: ICANN—I recall them being very specific. They did want something on a web page. It doesn't necessarily have to be WHOIS. It can be web access to RDAP.

ROB HALL: [I've] just always insisted at the registrar level that we have a website pointing to it because that's what the vast majority of registrants use. If

now we're moving that and we just have to run a RDAP service with no web interface to it, that's a great win for us. But I [inaudible]. I want to be careful they understand that.

RICK WILHELM: That's the gist of what we're going for on the registry side. That's what we're talking about, right? Because the whole point is that this decentralized notion of individuals either going to the registrar or the registry is not helpful to the people that are looking for this data because it encourages this decentralized view of it. They should be taking advantage of the bootstrap mechanism and going to lookup.ICANN.org. Right?

ROB HALL: Okay. If we can get that, that'd be amazing.

RICK WILHELM: Yeah. So I think that's part of the ... Yeah. Owen says, "Sounds good to me. I don't think ICANN would go for it." But that's part of what we're trying to do.

Brian has his hand in the air. Please, Brian, go ahead.

BRIAN KING: Thanks, Rick. Yeah, ICANN has really signaled that they're going to want registrars to continue to have a web-based lookup for RDAP. So essentially what Owen said in the chat. Thanks.

RICK WILHELM: Yeah. I wonder if—I'm going to call on Jim in a second—registrars would just [iframe] lookup.ICANN.org. I speculate about that idly, and then I'll turn it over to Jim. Please go ahead.

JIM GALVIN: I think that we need to stand firm as a group here on the CPH side, Rick, in the way that we've been approaching this. That is that there is not web-based lookup in RDAP. It simply does not make any sense. It is technologically the wrong answer. It does not carry forward to RDAP, and it should apply to registrars equally as it does to registries. We're pushing on that. We need to push together and just say, "Off the table." I realize that that's going to be a bit of a battle still, but I'm going to hold strong on that. That is a client responsibility. It is not the server's responsibility. We need to stick to that. It really is just that simple, as far as I'm concerned. We've had this discussion with him already—the pre-discussions—because we've got Francisco telling us one thing and Russ telling us something else and Russ representing something else. I'm sorry. I'll just be pretty blunt about it: Russ is wrong and we're right. And we should just stick to it, on behalf of both of us. Thanks.

RICK WILHELM: Okay. Very good. Any other ... Let me see. Brian and Jim, are those both old hands?

BRIAN KING: Mine is new, Rick.

RICK WILHELM: Okay. Brian, please.

BRIAN KING: Sure. I'm happy for us to agree on a unified approach, if that's what we want. I think we will be, just for transparency, battling years of learned behavior to go to registrar WHOIS. But I actually think Rubens is onto something in the chat: if we're slowly educating the community by including a link to the ICANN lookup, I think we could probably agree to host at least one link to an RDAP client, whether it's the registrar's own, if they feel like doing it, or the ICANN lookup page. That might help to change behavior that's been learned to go to the registrar. At the end of the day, if we are going to be aligned on this, let's take the opportunity and let's be aligned on it. I talked to our Mark Monitor engineers, and we don't really care. We don't think it's that big a lift to spit up an RDAP client. We're going to have an RDAP client anyway. It's built already. It wouldn't be that big of a lift, according to us. But, if not everybody wants to have to do that, I think we can align with the registries and push back on ICANN on this. Thanks.

RICK WILHELM: Very good. Thanks, Brian. Rob, new hand? Old hand?

ROB HALL: A new hand.

RICK WILHELM: Fire away.

ROB HALL: I think we want to be really careful, [if] just point them to lookup.ICANN.org, that we don't get into a situation where we can't offer our own client. So just be careful of saying, "Oh, well, just iframe it or point it there," because ICANN may grab onto that and say, "Great." And then we're stuck, where we have to do that and can't offer them because a lot of registrars love (and registries) to add advertising around these things and may see it as a profit center, not just a cost. So I think it should be optional for us.

I agree that we should try to hold a hardline to not have to do this as a mandatory practice.

RICK WILHELM: Very good. Thanks. Old hand/new hand, Brian?

BRIAN KING: It's old. Sorry.

RICK WILHELM: Okay. Very good. The one thing I'd say is that, right now, is remember, as we're talking about this web page, that that's part of our SLAs. So that's just something to remember as we think about it: in considering this, when we're talking about putting it up and keeping it up, it is part of your ICANN SLAs and such. So that's something to think about.

Let's move on to the next one here: RAA WHOIS accuracy program specification. The topic that we just passed there about the lookup page for both the registrars and the registries might be a good one to get an input on from the broader community of our respective stakeholder groups. That might be one where the broader group has some input.

So, RAA WHOIS accuracy program specification. This is: change references from WHOIS to registration data. There's no definition of the term "registration data." This term would need a definition/a main agreement. So this is some drafting work that needs to be done.

To me, this doesn't seem that bad. I'll defer to others. Maybe Brian can comment if this seems okay and reasonable. But this, to me, doesn't seem that bad to get done as I look at this section of the document.

So, unless there's any other comments on this one, we'll go down to the next section here: RAA registration data directory service WHOIS specification. Here we've got some, in the first section, definitional stuff that we're doing. Then here we are basically having some parallel updates to the RA document that is the definition of the profile and then an additional parallel update to talk about how the profile gets updated. So that should resonate with those of you who were paying attention during the first half of the discussion on the RA document. So that should be also similar. Then we've got the ramp-up period here also being defined towards the bottom of your screen—now in the middle of your screen; now at the top of your screen—as being the first 180 days. Then here we've got changes to the service levels as we scroll down. So here, on the service level, on the updates here, the only thing that's really material different is we've got the 5,000 milliseconds being

parallel to the 5,000 milliseconds. That is the same as the registry number. So you can see that this is different than the 4,000 milliseconds above. Again, flogging a horse on the 5,000 milliseconds, the rest of it is the same as before.

Kind of make sense? Any questions about the SLAs there that we've seen? Probably not, I'm guessing. I'll let people contemplate.

Then, right here at the bottom, this is the registrar version of avoiding double jeopardy—not being considered a breach of the RAA unless both RDAP and WHOIS have concurrent downtime, and each has exceeded the SLR thresholds. So this is preventing double jeopardy for the registrars. So this is also parallel to what we saw in the registries.

That gets us, according to my scrollbar, to the end of the document. We've gotten stem to stern here.

Now that we've gotten to the bottom, did anybody have anything that they expected to see or something else that was a little bit of a surprise in either the registry or the registrar sections that you'd like to go back and have us revisit?

Not seeing any hands. When I look back here, some change that we made, scrolling back up a little bit, is we added the web-based lookups in the registrars section to avoid that—in other words, get rid of the obligation for the registrars in that section right there, headed by Jim Galvin. Thank you, Jim, for doing the typing. Then we've got some stuff in the searchability in the registry section that we might want to do a little bit more editing on. But, in general, we made a little bit of progress there on that—some edits on that one—about how to tweak that. We

might want to spend a little bit more time on that one, maybe on the e-mail list or something like that. I think that that was the only other change that we made just in there on searchability.

We're going to have a little bit of time, so maybe we can spend a few more minutes and bolt down the searchability thing and just get it done-done.

But I just wanted to give a quick ... Let me see if Rob is still here. Rob, you had sent a note to the list earlier today. Do you think our discussion earlier at the beginning of the RA section covered your question that you had put out on the list? I wanted to make sure you just got airtime on that and—

ROB HALL:

I think it does. My concern was about two data sources and what if they're out of synch. And imagine what happens in a UDPR when someone points out there different. Which one is right? One person wants to use one for their advantage. I just really worry about not being able to proxy this.

The other problem is the jurisdictional one.

RICK WILHELM:

Yeah.

ROB HALL: And that's a bigger debate of who owns the data and who has the right to give it out? Can a registrar say, "In my jurisdiction, I can't give it out to the public, so, Mr. Registry, you can't"? So I'll give it to the registry for them, but they can't give it out because it's illegal in the jurisdiction where I am. That opens up a different can of worms, I suspect.

RICK WILHELM: Right. Sam, I see your hand up. Sam, please go ahead.

SAM DEMETRIOU: Thanks, guys. I'm not sure how helpful this is to addressing the concerns that Rob has raised, but I will point out—others who are also following this work can correct me or back me up on this—that I know that that kind of question about forum shopping when it comes to disclosure is something that the EDP Phase 2 Team has thought about. I think where the thinking is in this group at this point in time is that the queries by default would go to the registrar first, as the registrars are the authoritative source.

I see Brian is coming in, probably to rescue me from myself. But that's just something to keep in mind as we're grappling with this question, especially about the proxy. I think this is potentially an argument in favor of why the registrar should be reflecting its own data as the authoritative source for the registration data.

Brian, please go ahead and correct anything that I misspoke of.

BRIAN KING: I got your back, Jack. No, you're completely right. That is the idea. That received unanimity in the EPDP. [Baconator] has got your back, too, in the chat. So that's related to SSAD. I wouldn't be opposed to reflecting that in RDAP as well. I know that's already technically sorted out—how the bootstrapping works—which is more technical than I am. I just know the word “bootstrapping.” But I think that has something to do with it. But, yeah, gotcha. You're right.

RICK WILHELM: Very good. Rob, please go ahead.

ROB HALL: I just want to make sure that we're not ignoring this then, hoping that some other body in ICANN solves it. So I'm happy discussions are going on because it is a thorny issue.

Sam said something that I picked up on that I just want to make sure I head. You said the registrars are the first choice electronically in the RDAP protocol, or something along that vein. Did I mishear? I'm worried that somehow people are thinking there's going to be a hierarchy here and ICANN is going to pick one over the other first, based on what type of centralized WHOIS it is.

SAM DEMETRIOU: Sorry. That was imprecise phrasing. I think I should have maybe used the word “default,” that, in the context of the SSAD, especially in the hybrid model that's being considered now, where requests are centralized at one central gateway or portal but then they're handled by

the applicable contracted party, the default would be to go to the registrar record. That, I think, is what I meant.

But, again, guys—Beth or Brian—let me know if that’s not correct.

BRIAN KING:

That’s exactly correct. The only thing I would add is that there does remain the option, Rob, as you might be getting to, to do that ... I don’t want to call it form shopping because that’s not accurate, but there is the option in the SSAD—again, this is just SSAD context; this is not RDAP—to send the query to the registry. But the registrar is the default if the registry is not explicitly requested. Thanks.

ROB HALL:

So do we need to building into these contracts then exactly terms around that? Like, is the registrar the default in all cases? Is it only if there’s not a centralized registry? That worries me: someone else is deciding this and would not ask.

Secondly ... sorry, I’ve lost my second point. I’ll come back to it. Sorry.

RICK WILHELM:

No problem. Beth, please go ahead.

BETH BACON:

Thanks. I just wanted to again agree with Sam and Brian that, yes, that is accurate with regards to SSAD.

But also I think Sam was offering a helpful comparison, but I don't think she meant—and I don't think it would be accurate to say—that it's an apples-to-apple comparison between what we're talking about here and also what the SSAD is doing. The SSAD also has to take into account GDPR and the requirements that, if I as a registry have data and someone requested it of me, I'd have to give them what I have. But the thought, when we're talking about SSAD, is that registrars generally have the more authoritative source because they have the contract and the relationship with the customer.

So I don't think we need to mirror it, as, again, it's not, I don't think, an apples-to-apples comparison but just a guiding example. Thanks.

ROB HALL:

Should we be pushing to then ... Registry WHOIS came in—centralized registry—because no one liked the bifurcated model, where you had to go to one and then another. This RDAP protocol replaces that problem. Is there a reason for registries to give this out? We fought hard to go that way, but I'm just wondering what the new [technological advance] is. Is this really now just a [nice-to-have]?

BRIAN KING:

I can speak to that.

RICK WILHELM:

Fire away.

BRIAN KING:

Thanks. I want to be careful here. I think that the community spent a lot of time and energy on thick WHOIS, and I don't want to cause—pardon my French—a shitstorm if we try to impact that here. If that's somebody's goal or if that makes sense for us to do differently, let's talk about that.

I think, technically speaking, RDAP works how it works. I think what we're talking about here is our obligation to implement RDAP as it currently exists and as it currently works. Do we think we require any kind of changes to how RDAP works in order for the contracts to be amended? I don't think so. I know Rick and a lot of really smart people spent a lot of time figuring out how it works and making sure it works how it works. Thanks.

ROB HALL:

I have two quick points to that. One, I want to make sure that there's not values or behind-the-scene levers that ICANN no controls outside of this. When I heard the word "default," I wanted to make sure, "Wait, has there been a discussion around that?" because that may be something that we want to be put in these contracts: specifying who is the default, whether it be the registrar or the registry.

The second thing I want to say about this discussion is that we immediately went to, "Oh, yes, but it's going to be centralized because it comes from lookup.ICANN.org and then out to us." As I understand it, those are only queries that need authentication. I suspect the vast majority won't go that route. They'll come directly to us.

RICK WILHELM: Beth, please go ahead.

BETH BACON: Thanks. I don't want to beat this dead horse, but, again, it's not apples to apples. In the SSAD, Rob, the items that will be sent to the parties are the ones that need individual evaluation or response by the registry or registrar. The ones that actually don't need to be individually evaluated by a person and have a justification done are the ones that will be automatically available to authentication or verification.

But, again, I don't think that that's a decision for this policy to make—this particular group. It's certainly something we could discuss as a community, but I don't think that that's something that we want to start making contractual changes about. I think it would be a big deal. Again, it's not exactly the same concept. Thanks.

RICK WILHELM: Very good. Thanks. I just wanted to recognize Rubens' comment in chat about RDAP—not the first part of it, but just more that RDAP works on both by referring and following referrals. That's accurate. If you go to lookup.ICANN.org and you type in, for example, GoDaddy.com, you can get contact information for GoDaddy.com because it will chase the referrals as Rubens describes. So just to use that an example. So it works for both thick and thin.

Jim, please go ahead.

JIM GALVIN:

Thanks, Rick. I think that Rob is asking a question that is more applicable to client behavior. Then it's a question about what the policy is for a centralized query point.

Let me be a little more concrete about it. From the point of a registry or a registrar, if we get a query, we're simply going to respond to that query as is. I think that the question that you're asking really doesn't apply in that context. There's no default to this or default to that because, in this case—Rubens made this comment earlier in the chat room—we can't control client behavior. Clients are going to do what clients are going to do. They're going to query who they feel like querying. You're either going to have data for them or not. They might get lucky and ask you about something and you've got some amount of data and you're going to give it to the because that's what you do and that's what the system says and that's what all this contract language is all about.

On the other hand, what I think is the stuff that Beth and Brian and others are referring to is, when you get to talking about this centralized lookup system that ICANN has and the SSAD in particular—the authenticated system—it's going to be some amount of processing yet to be detailed and fully decided. Then it's going to have to do something to refer to the client to a better place for the data. At some point, that decision is going to happen, and that mechanism has to come into being.

There are a couple of ways that mechanism can happen, but one of them is, within the RDAP protocol, ICANN could tell the client, "Oh, by the way, go over here to get your answer," or the centralized server

could in fact itself go to some other place to get the data and then return it to the client. The policy that you're talking about is about responding to that decision point. That really has nothing to do with what's going on here. That has to do with that point in time. I think it's fair to suggest it's not in scope for what we're doing here. It is being very well-addressed in the context of the SSAD and the EPDP.

That's the way that I view that sort of issue. I don't know if that helps, and folks may disagree with me. But that's a little bit of something concrete there. Thanks.

RICK WILHELM: Okay. Very good. Thank you very much, Jim. Rob, is that a new hand or an old hand? Sorry.

ROB HALL: It is.

RICK WILHELM: Go ahead, please.

ROB HALL: I mostly agree, except that we have to keep in mind that the other party running this is ICANN. So they're talking about a centralized lookup point that acts as a client to us. So I don't know why we wouldn't want to be saying, "Hey, wait a minute. You're going to be running the largest, busiest client against us. Perhaps we should put something

about those rules in our contracts.” So you’re right. I am trying to say, “Wait a minute. This is about someone else’s client, but we have the power right now because they are the other party [inaudible] [conduct] with us.” That’s what makes me nervous: you’ve two parties of [conduct], and one is saying, “Oh, don’t worry. Trust us. We’ll figure this out,” and at the same time requiring steep SLAs from us.

RICK WILHELM:

Okay. Thanks, Rob. Anybody else care to comment?

Very good. Let’s see. We’ve got a couple minutes left. Did we want to try and look a little bit more at a few words of these comments on searchability? Or we want to talk a little bit about next steps? Those are a couple things we can do. Anybody have any thoughts about that?

Maybe we should – Jim Galvin, please go ahead.

JIM GALVIN:

Thanks. I think I’m a little nervous about next steps until we have a few other key people with us here talking about this, notably Donna and Jeff, of course. I think that we should make sure we have a more complete team, I guess, together before we make any firm decisions about next steps. So I guess that’s my real comments in all this.

I also think that we’ve made some comments in here. Certainly at least one next step is we have to resolve some comments so that we have a document that we can actually give to ICANN. What I understood as what we’re after here is we’re producing this because we want to give this to ICANN as a way to kick off negotiations at some level. Clearly, we

don't want to give them to it with any comments in it. So certainly at least one next step is resolving those, and having a plan for making that happen certainly seems prudent. Thanks.

RICK WILHELM: Yeah. I'd agree with that, generally. What's our schedule for our next meeting? Do we have anything booked right now, Sue?

SUE SCHULER: Same time next week.

RICK WILHELM: Same time next week. So why don't we do this? Why don't we plan on at least a session next week where hopefully we'll have a little bit more folks? We'll go through this. We'll work on resolving these comments and talk about our ... Rob asks, "Is next week with ICANN or with just us?"

I'm assuming it's going to be with just us. I think we said that we're not going to reach out to ICANN until we've actually gone to the stakeholder groups and then come back to ICANN. I think that's where are right now. If anybody has a different idea, please put your hand in the air.

So then I think what we'll do is we'll have a meeting next where we look to resolve the rest of these comments, clear up this stuff on searchability, and, probably more importantly, work on a, I would suggest, set of directed questions that we would take our respective

stakeholder groups. Maybe that would be a good set of goals for next week.

Does anybody have any further suggestions beyond that for next week? So a set of questions, resolve the comments, and tune the language on searchability. That will be our goals for next week. A set of questions for the stakeholder groups. Does that sound like a plan?

Seeing no—Brian says, “Good.” Catherine says, “Good.” Jim claps. He found the Clap button. I didn’t even know that existed. Jim is now a Zoom expert. He’s going to be trainer and maybe even train the trainer.

Okay. That sounds good. Anybody have any last-minute walk-ons that they’d care to bring up? Questions, comments, concerns?

Seeing none, I think, Sue, that you can take us out and wrap us up for the day.

SUE SCHULER: Thanks, Rick. Andrea, we can end the recording.

ANDREA GLANDON: Thank you. This concludes today’s conference. Please remember to disconnect all lines, and have a wonderful rest of your day.

[END OF TRANSCRIPTION]