
ICANN Transcription

New gTLD Subsequent Procedures Working Group

Tuesday, 28 April 2020 at 0300 UTC

Note: Although the transcription is largely accurate, in some cases it is incomplete or inaccurate due to inaudible passages or transcription errors. It is posted as an aid to understanding the proceedings at the meeting, but should not be treated as an authoritative record. Attendance and recordings are posted on the agenda wiki page: <https://community.icann.org/x/qi2JBw>

The recordings and transcriptions of the calls are posted on the GNSO Master Calendar page: <https://gns0.icann.org/en/group-activities/calendar>

JULIE BISLAND: Good morning, good afternoon, and good evening. Welcome to the New gTLD Subsequent Procedures Working Group call on Tuesday, the 28th of April, 2020.

In the interest of time, there will be no roll call. Attendance will be taken by the Zoom room. If you're only on the audio bridge, could you please let yourself be known now?

All right. Hearing no names, I would like to remind everyone to please state your name before speaking for transcription purposes and please keep phones and microphones on mute when not speaking to avoid any background noise.

With this, I'll turn it back over to Jeff Neuman. You can begin, Jeff.

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.

JEFF NEUMAN:

Thanks a lot, Julie. Welcome, everyone. Today's call will be technical in nature, but of course everyone should please feel free to ask questions. Some of this material today really gets into the weeds of some not-so-easy topics. So please do make sure you ask any questions if you have any.

The topics today will be internationalized domain names and security and stability—at least the portions of security and stability that do not deal with name collision. Name collision will be a topic for our next call. I'll talk a little bit about that after I ask if anyone has any updates to any statements of interest.

Okay. Not seeing anyone that's got any updates. That's good. Just a reminder: please do update any statements of interest as things change. I hope everyone is doing well, at least for me this evening, but for you tomorrow morning or wherever you are in the world.

As I was saying, as today's topic we're going to talk about internationalized domain names. There's some very specific, very technically-detailed proposals that we're making, many of which were either recommendations or suggestions from the SSAC or from the already-completed policy processes, although they were not GNSO—or some of them were, but some of them were also committee working groups either from the Board or from just various working groups set up in the community to deal with internationalized domain names over the years. So we'll get into that in more detail, and then we'll talk about, like I said, security

and stability, mostly focusing on issues around root zone delegation and the rate of delegations.

On Thursday ... Yeah, that's right. I'm trying to think of the days because it's hard to keep track of days. So our next call this week, on Thursday, will be to talk specifically about name collision. For that, I've invited the Name Collision Analysis Project Co-Chairs—there are three of them—to the call on Thursday, and I've circulated a draft final report that was submitted by the independent consultant that was hired by ICANN, specifically the OCTO organizational part of ICANN. In that, you will notice what was added to the original report, for which the comment period recently completed, was the consultant's recommendation on whether to conduct Studies 2 and 3. Please do read that. It is still very much in draft form. The Name Collision Analysis Project discussion group is still discussing it, so that report is very much subject to change, but still it'll help you prepare for the call on Thursday.

With that, also another reminder that, next week, on May 4th, will be the little-bit-longer call. So it'll be a two-hour call instead of an hour-and-a-half. On that one, we'll talk about primarily auctions/mechanisms of last resort and private resolution of contention sets. So we put that as a good, longer discussion item. So that, like I said, will be May 4th.

On May 7th, we're going to talk about predictability as the main item. Just a reminder that, if you have any proposals on what we were talking about a couple weeks ago on GAC advice early warnings, please do have those submitted by May 1st so that we can see it on the list and then discuss it on May 7th. If there are no

proposed wording changes, then we'll take that subject off the May 7th schedule and just talk about the predictability framework.

All right. Sorry to spend so much time on admin, but it's good to get that done upfront so that we can spend the rest of the call on substance.

Let me just stop there and ask if there are any questions while Steve brings up the section on internationalized domain names. And Steve did that before I even finished the sentence. Thanks, Steve.

Okay. No questions. Good. We'll move on to Section 2.7.5: internationalized domain names. We lead off that section with an affirmation. Actually, it's with a slight modification because the principle from 2007 just talked about that 2012 round as opposed to the ongoing commitment to internationalized domain names. So we're proposing a modification so that it states that, "New generic top-level domains should continue to be an integral part of the New gTLD Program." The original principle just had it as a one-time deal as opposed to being an ongoing principle.

Thanks, Steve. Steve has posted a link to the Google Doc that we are working off if you want to follow along in the Google Doc.

The first recommendation that we have in here, which is labeled "Recommendation with Rationale #2," states, "Compliance with root zone label generation rules must be required for the generation of IDN TLDs and variance labels."

Then there is some bracketed language that was proposed to be added, I believe, from the ICANN Org comment. If we add that

comment—we'll discuss that in a minute—it would state, at the end, "including the determination of whether that label is blocked or whether it's allocatable," which I guess is word. Essentially, it's asking whether the variants can be allocated. We'll talk about who it could be allocated to in one of the subsequent recommendations below.

Any questions on that? Specifically on the added language.

Okay. No one has got their hand up in the chat. Rubens says, "The added language sounds okay to me." Good. Anybody else have any comments on it?

Jim states, "'Allocatable' is in fact a word." That's good. I figured, since it didn't show on a spellcheck, it probably was.

Okay. Obviously, we're not making definitive decisions on this call, so we will include that text without the brackets. Of course, others can offer comments later on if they would like.

Under this recommendation, there are two implementation guidance or ... What's the plural of "implementation guidance?" Jim, I'm going to ask you that question. So you can look that up. Would it be "implementation guidances?" I don't know. Anyway, there's two of them. The first one states, "To the extent possible, compliance with IDNA 2008." And there's, "RFC 5890 to 5895, or its successors. and applicable root zone label generation rules should be automated for future applicants. The working group recognizes that some instances of manual analysis may be required."

The next one states, "If a script is not yet integrated into the root zone label generation rules, applicants should be able to apply for a string in that script, but it should not be delegated."

Does that make sense to everyone?

Well, let me ask a question. When I read this statement, it made sense that they could apply for it, but then you say it should not be delegated. So are we saying that it should go through then entire application/evaluation/objection processes and then wait until the script is integrated into the root zone label generation rules, or should we just hold the application until that integration happens and then, at that point, you start the evaluation, the objections, and all the other stuff?

Jim is providing the plural version. Thanks.

Rubens is saying, "I think this is weird. The application should be on hold and not go to contracting." So, Rubens, what you're saying is that you think it should go all the way up until the approval but should not go to contracting until the label generation rules for that script are integrated?

"Yeah." Okay. Paul, go ahead.

PAUL MCGRADY:

Thanks. Super-dumb question. Is there any possibility that one application that is held up at the beginning could be so similar to another application in a script that is already in that it could cause a contention set and so holding up an application at the beginning rather than at the end would hold up another application?

Because, if that could happen, then we probably want it to be held up right before delegation rather than at the beginning. Thanks.

JEFF NEUMAN:

Thanks, Paul. We're talking about variants at this point, aren't we? No, I'm sorry. We're talking about the root zone generation rules. I think it ... That's a good question. It would be a very, very, very much edge case.

Rubens is saying, "I believe the scripts where such similarity could happen already have label generation rules, like Latin and Cyrillic."

Yeah. Many scripts—I shouldn't say "most" because there's probably hundreds if not thousands of scripts around the world that we have no clue of ... I think it's very unlikely.

I see Karen has pointed to a list of all the root zone generation rules for all the scripts.

I think what we could say is it would still be in line with what Rubens had proposed, which is that you take it all the way through, up until contracting, and then you hold then contract until the label generation rules for that script were finalized. So, if there was a contention set, that would have already been processed and bringing it all the way through. That seems like it would cover that in the very rare case that that could happen.

I saw Jim's hand up but then the hand went down. Jim, is that what you were going to say?

Okay. I do think what we should say at the end of this “but it should not be delegated”—I don’t want to the exact wording on the fly here—that—is, basically that those applications should go all the way through the application evaluation and objection processes, up until but not including the contract phase, so that you could take care of any of the other issues. And the only thing it’s on hold for is the contract, which of course is pre-delegation anyway.

Justine states, “Along the lines of thought of Rubens and Maxim, are we giving applicants sufficient warning that contracting may be held up subject to root zone label generation rules?”

Well, most scripts have root zone label generation rules for the majority, if not all, of the scripts that, in theory, really would be applied for. So applicants will be aware of which scripts have label generation rules at the outset. So I think what, Justine, you’re saying is that we should also make it clear that, if they are applying for a script that does not have label generation rules to determine variants, then they’ll have the warning that contracting and ultimately delegation may be held up.

Maxim states that there should be an explicit warning. “I think it was around one or two years.”

I think that’s right, Maxim. There should be a warning. I think the one or two years was specifically because we were just starting to write those tables. I think, by now, hopefully the process would not necessarily take so long, since we’ve done a number of them.

Jim, go ahead.

JIM PRENDERGAST: Thanks, Jim. Not necessarily on this particular section, but I was reading through Package 2 that was sent out last week. I know, in the testing section, there is reference to be adding [IDN] tables. So I just want to call out that there should be some synchronization between what we talk about in Section 2.11.1 and here. It should be consistent in regards to the tables that are necessary for passing either pre-delegation testing or the advanced testing for registry service providers. Thanks.

JEFF NEUMAN: Thanks, Jim. I think that's a good suggestion. We definitely should refer back to the registry system testing, where we do talk about those. Or is the evaluation or both? I'm trying to remember now which 2.11 ... You said it and I totally blanked. But, yeah, I do think that there should be a reference back.

Steve, go ahead.

STEVE CHAN: Thanks, Jeff. I highlighted a comment that we put in in the drafting of this section. There was actually a recommendation in this section about the IDN tables, but, to Jim's point, they seemed duplicative. So the point of this comment here is to inform the working group that, for that purpose, we actually did not include it here, but we wanted to make sure that you see what the recommendation in the registry system testing section looks like. You can see it in the text that's on the screen here. If there is desire to include it in two spots, I suppose that's okay, but the

other possibility is that we could just make sure that the cross-reference is included in Section D, which I just took a look at. I forget to put it there. So there's a couple of options. You could have it in two spots or you could just make sure there's a cross-reference in Section D to make sure it's accounted for. Thanks.

JEFF NEUMAN:

Thanks, Steve. I think we've been trying to avoid having things in two spots in case wording gets changed on one of them and, for some reason, it's not consistent or we forget to do the other one. So my personal preference—of course I'd love to hear comments from the group—is that we not state the same thing twice.

I'm just reading then chat. Rubens states that IDN TLDs is a totally different subject from IDN tables. Justine says, “[But what Rubens says is also important:] top-level versus second-level.” Then Rubens states, “IDN tables apply to all TLDs.” “Right, including ASCII non-IDN TLDs as long as the registry offers IDN registration.” Yes.

So, Rubens, are you saying here that the wording should change? Because here we're talking about the IDNs at the top level, I think.

“No wording change. Just to remind to not include.” Okay. So do not include the IDN tables here. Got it.

Looking at Rationale 3, this was a subject that we had talked about and actually was heavily discussed before the 2012 round began. But, because the subject was introduced so close to the start of that 2012 round, it was not allowed. This is the topic of the one-Unicode-character gTLDs, which, of course, even though it's

a one-character internationalized domain string, the reality is it's always going to be XN--. So it's always going to be, in reality, more than one character. It's going to be at least five.

Anyway, these gTLDs may be allowed for limited script language combinations where a character is an ideograph or an ideogram and do not introduce confusion risks but rise above commonplace similarities consistent with SSAC and joint ccNSO-GNSO working group reports. So here we're just taking the recommendations from SSAC 052 and the—I'm trying to remember—joint implementation group of the ccNSO-GNSO Working Group; their two reports.

This next one was originally a two-parter that we've put into one part because one of the parts seemed very redundant with a recommendation that was a couple above this one. Now it states, "IDN gTLDs deemed to be variants of already existing or applied-for TLDs would be allowed, provided they have the same registry operator"—we'll get to the bracketed language in a second—"implemented by [,by] force of written agreement, a policy of cross-variant TLD bundling."

So this essentially, in layman's terms, means that the same registry operator must control both variants and that you cannot give one variant to one registry operator and another variant to another registry operator; that variants must be allocated to the same registry operator because of what we're talking about here of bundling rules or other consistent policies.

Justine is asking if we can just say that in layman's terms.

I think the last part is a little ... I think we might have copied the wording from an SSAC recommendation or from the joint implementation group. So I would personally like it stated in more layman's terms, but, if it's quoted, we probably should quote it. Certainly, we can do it in the rationale if we can't change the wording here.

What I do also want to talk about, though, is the added language that was suggested by, I believe, the Registry Stakeholder Group, which not only wanted the same registry operator but also to mandate that the same backend registry services provider has to also be committed to in the variant applications and ultimately the variant TLD registry agreements. This is an interesting one because backend registry service providers do not have direct contracts with ICANN. It obviously makes a lot of sense in terms of that I'm almost 100% positive that a registry operator would have to use the same backend provider to operate variants for purposes of bundling and/or allocation just by feasibility. But, to me, requiring it as a hard recommendation seems a little bit out of place, simply because there's no privity between ICANN and the backend operator.

Rubens says in the chat that implementing it with backend RSPs would be a nightmare," which I agree with, but I'm just looking at it in terms of who ICANN has got the relationship with. If a registry operator figures out some way to do it with different backend providers, I'm not sure we need to prevent that.

Alan, go ahead.

ALAN GREENBERG: Thank you. Jeff, I was going to say something different, but, in response to what you just said, you're implying that, if the registry operator is innovative enough to figure out how they could run the two different variants on two different backend providers, we should allow them. I think that's what you just said. That may well be true, but we're trying to build a stable, reliable system, and the registry operators belief that it will be okay, I think, pushes credibility too much.

So I like the rule as written, and I don't see a contractual problem with it. We're contracting with the registry operator, and the contract says, "You cannot run two different variants on different backend providers." Now, it's the registry operator who has the obligation to make sure they're being run consistently. The simple way of doing that, among other things, is to say they have to be on the same backend.

So I don't see a contractual problem with it. I think we're on ground that's a little bit shaky to begin with because people have always had some qualms about variants. I think we need to be safer rather than flexible here. Thank you.

JEFF NEUMAN: Thanks, Alan.

Maxim states, "But if you can properly sync, why not?" Rubens says, "Maxim, it's the same [raised] condition that prevents a string from having more than one registry system at a time."

Anybody else? I guess my thinking was we're already saying the same registry operator. Are we being too prescriptive for ... I can't

see the actual benefit of being this prescriptive. The point is that you're having the same registry operator and the same rules, and the registry operator is the one that sets the policy and the rules, not the backend provider. I just don't think we should be overly prescriptive if there are other ways to achieve the same thing. But that's just a personal view. I'm not sure there is a way to do it, but I just don't think we should be prescriptive for the sake of being prescriptive.

Karen, go ahead.

KAREN LENTZ:

Thank you, Jeff. I was looking for the reference, but I wanted to pointed out that the recommendations on IDN variant TLDs were passed/approved by the Board, I think, last year and then referred to the GNSO and ccNSO for their consideration in policy work. It actually does address this point. I'll put the link in the chat. It says, "For feasible and consistent implementation of these requirements, the same backend registry service provider, if applicable, must be employed for operating all of the activated IDN variant TLDs by the registry operator. I'll put the reference in the chat.

JEFF NEUMAN:

Okay, good. Well, that seems to answer it then. Maybe that's why it was bracketed.

Does anyone then have a problem? I'd like to cite that. Karen put the link in there. If we can cite that report as to the source of that recommendation, I think that that would be better.

Alan, go ahead.

ALAN GREENBERG: Thank you. What I was going to say before Karen did her intervention is, aren't we looking at a real, real edge case? Are there many TLD operators/registries that use multiple backends? I was going to ask, is this something common enough that we have to consider it and then worry about that we're imposing too many rules on them? It sounds like an edge-of-an-edge case to me.

In any case, it seems to be settled, so so be it.

JEFF NEUMAN: Yeah. Thanks. There are a few registry operators that use multiple backend providers, so it does happen where there are multiple TLDs and, for whatever reason, a registry operator wants diversity of service providers.

Since we found—Karen, thank you for that—there is support from that particular ... We took it from a recommendation that exists, and it was adopted by the Board. I think we should keep that in there and maybe even quote it as opposed to ... I can't remember if we rewrote this for any reason, but, if it works better as a direct quote, then perhaps we should do that.

Karen, your hand is still up, but I think that might be ... okay.

The next one is ... Here's where we get really in the weeds, and I had to read this so many times before I think I understood it. It says, "A given second-level label under any allocated IDN variant

TLD must only be allocated to the same entity/registrant or else withheld for possible allocation only to that entity.” Actually, this one I got. Then next one was harder. In other words, if there is a registry that does allocated variants as second-level strings, then they have to have a policy in place to make sure that the same registrant has the string and all of the variants either blocked or just allocatable to—I guess that’s a word now—the same entity. That one seems pretty straightforward.

It’s this next one which ... Well, I’ll read it. “The second-level IDN variant labels that arise from a registration based on a second-level IDN table—all allocable IDN variant labels in the set—must only be allocated to the same entity or withheld for possible allocation only to that entity under all allocated variant TLD labels.” I know that’s a difficult one because there’s several sublayers, but it’s the same principle, essentially, as the previous recommendation that variants within a TLD, as well as within variant TLDs, are allocated to the same entity. So, if you have two TLDs at the top level that are variants of each other, then the variants at the second level of the first TLD and the second TLD must be given to the same registrant.

Does that make sense? Basically, it’s the transitive property in math that says, if A implies B, and B implies C, then A implies C. In other words, if you have two IDN top-level domains that are considered variants of each other and, in both those top-level domains, the registry operator decides that it’s also going to allocate variant second-level names, then the variants in the first top-level domain are by definition going to be variants of the

variant second-level string in the second TLD and, therefore, all of them have to be the same entity.

Paul, does that make sense, if you rewrite it a mathematical equation?

PAUL MCGRADY:

Yeah. Thanks. I get that. I have an even dumber question than the one you answered. You gave me a lot of credit by giving me a smart-person answer, but I'm going to [need a further] dumb-person answer, which is: because people using those second-level domains will bump into each other, people will be confused getting there. It will end the universe as we know it. So, I guess, what's the harm?

JEFF NEUMAN:

It all deals with user confusion. Right now—well, I shouldn't say all, but all the ones that at least I'm familiar with—all of the registries bundle second-level variants. In other words, if you register one second-level string, and that second-level string has a variant, most registries will bundle them to either block one of them or give ownership of the two different variants—assuming there are only two—to the same entity. That's because they look identical to each other and, therefore, even though you could technically use them both, it would confuse the hell out of the end users because they look exactly the same.

ALAN GREENBERG: Jeff, you may want to defer this discussion until we look at the next recommendation because it becomes much more interesting.

JEFF NEUMAN: Hah-hah. Okay. Well, let me just continue on this. Paul, if it's extremely confusing to end users at the second level, now you add the complexity of having two top-level domains that are variants of each other both allowing second-level registrations of variants and you have—let's say there were two variants of that second-level string in the first TLD—the same at-least-two variants in the second-top level domain. Therefore, you now have four variants of each other that all need to go to the same entity.

Following this discussion to the next one, which Alan rightly states is even more interesting, second-level labels derived from the above recommendation are not required to act, behave, or be perceived as identical. I can't remember where we got this one.

Let's go to the rationale, as Steve says. So we're going to Rationale 5, which explains the two recommendations, and then I'll go to Alan. "For similar reasons as indicated in Rationale 4," which is, "in support of security and stability and in light of the fact that IDN"—I'm reading Rationale 4—"variants are considered to essentially be identical, the working group believes that IDN variant TLDs must be owned and operated by the same registry operator. To the extent that the TLD were to change hands at any point after delegation, IDN variant TLDs must remain bundled together. Accordingly, IDN variant TLDs should be linked contractually."

Now we go to Rationale 5, which states, “The working group believes that second-level IDN variants should only be allocated or reserved for allocation to the same registrant.” This applies both when it is a certain second-level label under multiple variant IDN tables and variants at the second level derived from the registry operator’s approved IDN table. So, even if all of the variants are not necessarily allocated according to a particular table, the registry operator needs to apply the same rules to all of the potential variants, whether or not they’re allocated.

“However, the working group, taking note of public comments received from the SSAC, agrees that second-level variants should not be required to behave exactly the same. Assuring that second-level domains have behaved the same has not been found to be technically feasible in the DNS. In addition, there are practical reasons for second-level variants for not be the same. Example: Simplified and traditional Chinese second-level variants could have the content on the respective webpages available in simplified or traditional Chinese consistent with the DNS label.” So, even though you’re giving the two variants—let’s say, in this case, it’s the traditional and the simplified Chinese IDN second-level variant—to the same entity, that registry operator can point one of the variants to a content that contains simplified Chinese language content, and it could also allocate one of the variants to the traditional Chinese content website.

Maxim states that this one is troublesome. Maxim, can you help with understanding that? Then I’ll go to Alan. Maxim, is there a way you can explain that comment?

MAXIM ALZOBA:

I wonder how we link contractually two—hopefully three—contracts with are independent now because, currently, a registry is a holder of the particular contract. It might cause issues later. For example, if there are three IDN variant TLDs and, for some reason, one of those of is not used by a registrant, and a registry decided to basically terminate one of those because of no registrations, for example, should it mean that a registry doesn't allow it to take down one or to kill other successful TLDs because one is not used? It needs to be clearer because contracts should be separable. If there is a policy requirement that says it should be [single], we don't need to things like, "Please write it into one single contract," because one single contract means one registry fee, and it makes a situation where some particular entity just pays less because they bundle. [It's wrong]. Thanks.

JEFF NEUMAN:

Thanks, Maxim. I don't think it's an insurmountable problem. As we talked about, everyone gets the same base registry agreement, but, in the case of IDN variant TLDs, there's no reason why you couldn't do a new specification that talks about how you handle the multiple variant TLDs and provide for things like the fees. So you could still have one contract, but, in that separate specification, you have those provisions that would only be related to registry operators that operate both a string and variants. So there are ways that, I believe, the legal team could certainly do that, but I don't know if we should be prescriptive on figuring out ways that they could do that.

Let me go to ... Alan, let me just finish this up with the comments. It says, "Jeff, could you retire a variant, or are you stuck with it forever?"

I think that would be mentioned in Specification 14, or whatever the next one is—15?—that could state that ... I don't see how ICANN would force or could force any registry to keep all of their TLDs.

Let me just read Rubens. That says, "Paul, that's probably something that a registry operator could specify in a RCEP request: "Stop Variant X from being offered."

Yeah, Rubens, there's multiple ways, just like how, now, any registry can just voluntarily give up their registry. Of course, they'll have to submit a transition plan and all that other stuff. Of course, then ICANN would be restricted because of these variant rules. They couldn't allocate it to anyone else. So there's that.

Karen states, "@Paul, the variants do not all have to be active. However"—right—"if TLDs are delegated as variant labels based on label generation rules, their status as variants would not change." Right. That's right, Karen. So, whether one uses it or not, they are always going to be variants of each other.

Maxim is saying, "We do not need to require killing all variants instead of retiring a single variant of the bunch." Right.

Justine is saying, "Might we need another implementation guidance regarding another specification?"

Well, we could, but, again, I just suggested one possible way that it could be done. There are probably certainly other ways it could be done, and I don't know if we need to be so prescriptive here if ICANN Legal or their advisors think that there's a better way to do it. So the answer is: we could certainly have implementation guidance. The question of whether we should, however, is the same question of whether we need to be very prescriptive. I don't know the answer to that.

Alan has been waiting very patiently, so, Alan, please go ahead.

ALAN GREENBERG:

Thank you. I now have about seven different things to say based on all the conversations that have gone before.

In terms of the contracts and the retiring of registries and stuff, I think the lawyers are going to have an interesting time ahead of them. They're going to have to either decide, as I think you implied, Jeff, to come up with a single contract that applies to two TLDs at the same time, or you need two different contracts that somehow cross-reference each other and are tied together. Either of those, I think, is going to be an interesting challenge. The concept of retiring one of the variants, if there are second-level domains allocated to it, might be really interesting because, as you point out, if you as a registry operator decide you don't want it anymore, ICANN normally will try to find someone else to operate. But that's not possible here. So there's some interesting challenges ahead in covering all of these possibilities.

I originally put up my hand to talk about the third recommendation that says they don't have to behave the same. That's a really interesting one because, as you point out, you may well want them to not have the same contract. You may want one in simplified and one in traditional or something else like that. The first two rules say the registry cannot be complicit in having the second-level domains act differently because they have to give them to the same registrant. It's an interesting challenge, by the way, of: can those two domains be done by two different registrars? But that's a registry problem, I think, not our problem.

However, it also does say, since ICANN cannot look at content, although the registry cannot be complicit in having the domains act differently, the registrant can. In fact, they could be different not only because one is in traditional and one is in simplified Chinese but that they could be completely different websites. They could have no relation to each other, even though they have effectively the same name. That's really outside of ICANN's remit. I don't think we have any recourse to try to fix that problem. The hope, presumably, is that they act in a rational way based on [there] being the same domain, but I think it's completely out of our control. It's sad, but I think that recommendation is the only thing we can say. Thank you.

JEFF NEUMAN:

Thanks, Alan. I do know of one case where a registrant wanted two variants and wanted to activate both of them to point to different places because of the traditional and simplified Chinese issue. Actually, I know of two cases, one where there was a

variant with a Japanese and a Chinese label and they wanted them to go to different places. So it can happen.

So I agree with you, Alan, that this is probably the most that we can say anyway, but I'm not sure that that's a bad outcome. All you're saying is that the registry has to make sure that the same registrant has the variants. How the registrant chooses to use it? So long as it's in compliance with the overall registrant policies. In other words, a registry may choose to have rules not to allow variants to be used, or a registry may have rules where they can both be used. But, subject to that, we can't control what the registrant does if both are allocatable.

ALAN GREENBERG:

Yeah. May I have a follow-on? We cannot control that. The registry can. What you're talking about is that they point to two different sites, but they both presumably have similar intent. One isn't a clone of eBay and the other a porn site. The registry, should it choose, could put rules on the registrant, saying that the variants that are allocated at the second level have to be similar or related or whatever—I don't know what the right words are—websites. But there's certainly nothing we can say.

JEFF NEUMAN:

Yeah. That would be such an edge case because a registrant that is sophisticated enough to operate two variant second-level strings and be smart enough to design content aimed at two different audiences to understand? It would be a such an edge case. But, yes, you're right. We don't regulate—

ALAN GREENBERG: Yeah. But it's out of our hands anyway.

JEFF NEUMAN: Right.

Paul asks a question about transfers. Again, you can't just transfer one variant. They're all bundled together in some sort of way.

I see Rubens and others have also responded to that.

ALAN GREENBERG: Jeff, do we need a recommendation that the UDRP and URS be modified to handle these variants, that multiple second-level registrations have to be handled as a group?

JEFF NEUMAN: Hah-hah. It's a good one. Very, very edge. The reason I'm laughing is because it came up in one case. But I don't know if we know enough about that issue to make any recommendations.

Let me go to Paul and Maxim. Paul, go ahead.

PAUL MCGRADY: Sorry about that, Jeff. Didn't mean to cut you off. I don't know that we need to do a recommendation here, but maybe we write a nice letter to Phase 2 of the RPMs PDP, alert them of the issue, and ask them to consider it.

JEFF NEUMAN: Okay, yeah. Then I'll—well, I'll go to Maxim—tell you about the one situation that came up, just to give an example. Maxim, go ahead.

MAXIM ALZOBA: Potentially, there could be situations where ... Since we are talking about [muting] these, it's not very dangerous, but speaking ... If we're limited here to new TLDs, then handling these groups in bundles might be simple but not necessary, for example, in a case where, for some reason, one entity has trademark rights for the word in simplified Chinese and in some particular class of services, and the other entity has the right for the other variant. It might be troublesome. Thanks.

STEVE CHAN: Jeff, if you're speaking, I think you might be on mute. Or maybe my line—

JEFF NEUMAN: I am on mute. Thank you.

STEVE CHAN: Hey, there he is.

JEFF NEUMAN:

You're right. I was on mute. So it's an interesting issue because it is ... I'll tell you the situation that came up in one of the cases I'm aware of. A particular trademark owner wanted to register their trademark in a TLD. It was an IDN string. But they couldn't because that string was considered a variant of a string that was registered by someone else. It turned out that the someone else that it was registered to was using it in an infringing manner.

So the party that wanted this string—the trademark owner—filed a UDRP and won the case, and so the UDRP ordered that the variant be transferred to the trademark owner. So the trademark owner now had owned the variant[s] of the string that it wanted, but the trademark owner couldn't convert that variant over to the string that it wanted because of the fact that registries don't have policies where they allow variants to be substituted.

So the only option was for the trademark owner was to either give up using that string and just deal with it and use the variant or let the string delete and take their chances on reregistering the other variant version of the one it wanted.

Hopefully, that made sense. I know it's a little complicated. But I only heard of that happening once. So it's not really a UDRP issue. It's an issue of a registrant being able to substitute the use of one variant for the other. It just so happened to win the name through the UDRP. Anyway, I don't know if I lost you all or made any sense.

Anyway, I'm reading Rubens' comment. "UDRP is meant to address a large share of trademark issues, but there's always the possibility of only a court being able to settle a dispute. As long as

the UDRP keeps working for a near totality, it doesn't have to deal with all possible cases."

Yeah. I think this is such an edge case, but, except for the fact that if a string is awarded via a UDRP, just like a string is transferred in any other way, all of the variants have to along with that second-level string if a transfer is to happen. But I don't think that's an issue for new TLDs. That's probably more of an issue of ongoing operations in both the transfer policy as well as the UDRP or any other action which requires the transfer of a string.

Let's go on then to the second subject that we're meant to talk about today, which is security and stability. We're really going to focus on issues involving the rate of delegation of [g]TLD strings.

We affirmed Principle A from 2007m which states that new top-level domains must be introduced in an orderly, timely, and predictable way. We also affirmed Recommendation 4, which states that strings must not cause any technical instability, and we add the following recommendations. One, ICANN must honor and review the principle of conservatism when adding new gTLDs to the root. That comes directly from one of the SSAC papers. It's probably the same one that we sight a little bit later on.

To give a little background before the next recommendation, if you recall the, the rule was set for the 2012 round that no more than one 1,000 delegations per year could take place. There's a lot of history behind that. We did a couple comment periods on this. Essentially, that was not what was mandated by the SSAC, nor was it what was mandated by the root server operators. The SSAC and root server operators and other technical groups

advised ICANN to be conservative and to focus more on the rate of change to the root zone as opposed to a fixed number of adding TLDs to the root.

As a result, and in accordance with those papers and the comments that we got, the next recommendation states that ICANN must focus on the rate of change for the root zone over smaller periods of time—for example, monthly—rather than the total number of delegated strings for a given calendar year.

Then what we state in it as implementation guidance is: the number of TLDs delegated in the root should not increase by more than approximately 5% per month, with the understanding that there may be minor variations from time to time. We did a whole chart on this. I can't remember if we referenced the chart here, but we show the curve of what would happen if the rate of change were increased by no more than 5% a month. What you get is that maybe, in the first year, it's closer to 1,000 delegations, but, in the second year, it goes up and allows for the rate of change to be that which is recommended by the comments that we got, both from the SSAC as well as the root server operators.

There are some questions. Let me go to Maxim and Alan.

MAXIM ALZOBA:

Actually, as I understand, there was no basis for 1,000. It's just a historically used number.

Speaking about that 5%, the justification in the paper from SSAC, I think, was just not very mathematical. Thanks. The particular

number—5% or 6%--wasn't mathematically justified based on experience of sorts It was quite weak, I'd say. Thanks.

JEFF NEUMAN: Thanks, Maxim. As we went back in the history, it really was a combination of the conservatism principle along with discussing the evaluation ability and the implementation of delegations by ICANN staff, where they came up with the 1,000 number. So it wasn't a purely technical guidance. It was "Here's the principle of conservatism," along with ICANN staff at the time saying, "Hey, look. We think we can process about 1,000 delegations in a year." So that probably is as good a number as any. So it's not that there wasn't any justification for it, but it was a couple principles put together as opposed to some neat mathematical formula.

Alan, go ahead.

ALAN GREENBERG: Do we have any statistics? I presume they're achievable, but I haven't seen, nor have I looked, for what the actual delegation rates were, either as a percentage or numbers over any period—per month, per year. How close did we come to the 1,000?

JEFF NEUMAN: We didn't come very close to the 1,000 at all.

JEFF NEUMAN: I know we didn't do 1,000 in a year. That would have been just ludicrous, given how we know the world unfolded. I'm just

wondering what the actual numbers were. ICANN was working as fast as it could to get these done. I'm wondering just what were the numbers like. It's perhaps not important today, but we're supposed to be doing [stat]-based policy, and I don't know what the [stats] there are at all.

JEFF NEUMAN:

Those stats do exist. There was a whole root server operator ... There were several reports that were done that talked about the rate of delegation being well below what the root could handle. So what was clear is that the numbers of the actual delegations didn't move the needle at all in terms of what it could handle. So that's really all we have. Hopefully, as ICANN does this "more" more, economies of scale and processes will become much more efficient, where it'd really become very much more likely that ICANN would be in a position to do more than 1,000 a year. If some of the wacky projections come to fruition where, let's say, you do have 10,000 applications—I'm by no means suggesting that that's what's going to happen—then you don't want to be telling the applicants, "Okay. Well, some of you are going to have to wait ten years before your delegation." It's just not a good outcome.

Steve posts in there that, yes, the 1,000 was more of an operational assessment into the security and stability recommendations.

Let me go to Karen.

KAREN LENTZ: Thank you, Jeff. Responding to Alan’s question about statistics, I’m pretty sure we presented this before and have that data. I’ve been trying to find it. But rather than trying to paste something in, we can take an action to follow up with the data. Thanks.

JEFF NEUMAN: Thanks, Karen. I think it might actually even be in that report that we were talking about. I’m trying to remember who did it now. It was several years ago. But, yeah, we certainly have those stats.

Maxim is saying, “I’m not sure the 10,000 is real now.” That may be. So this may all be for nothing, but, at the end of the day, if it does come to fruition, we don’t want to be telling people that, because of this of this artificial constraint, people are going to have to wait ten years in order for their TLDs to be delegated. So the recommendations here are in line with the reports and the recommendations that we got from the SSAC, from the technical community. So we’re not just making this up.

The next one is that ICANN should structure it’s obligations to new gTLDs so that it can delay their addition to the root zone in case of DNS service instabilities. Objective criteria should be developed to determine what could be classified as a service instability. The first part of this—the first sentence—came directly from one of the SSAC papers. The second sentence we added because, for the principle of predictability, it would be good to know what service instability actually means. So the exact wording of DNS service instabilities, I believe, came from the SSAC itself. So we’re adopting it, but we’re saying that you probably develop criteria to

determine what is a service instability. Again, it has not happened before, so we don't know how likely that is.

The next ... In fact, I should go back and see all of these implementation guidance are taken from the SSAC comments that were submitted to us in response to the initial report, and all of those were taken from previous SSAC papers, which are cross-referenced in those comments.

The next one is that ICANN should investigate and catalogue the long-term obligations for root zone operators of maintaining a larger root zone.

The next one is that the Office of the Chief Technology Officer (OCTO) should consult with PPI, Verisign, and the root operators via RSSAC, and the larger DNS technical community, on the implementation of these recommendations. That also came from the SSAC. The SSAC did cite the Officer of the Chief Technology Officer and these groups. When I first looked at this implementation guidance, I thought, "Well, why do we need to be so specific as to say the Office of the Chief Technology Officer?" because who knows what that might be called one day? But that is the wording that the SSAC used, so why change wording?

The next one is that ICANN should continue developing the monitoring and early-warning capability with respect to root zone scaling. That also came from the SSAC.

So those are all grouped together. Any questions on those before we go to the next recommendation?

Rubens actually makes a good point that we should replace “Verisign” with “root zone manager,” even though Verisign is the root zone manager. It should be the role as opposed to the organization that we point out. I think that is a good change.

Anyone object to that?

Okay. That seems to be supported. Good.

Recommendation 2 is talking about emojis and not allowing them. This was the subject of SSAC ... I don’t remember which one but it’s in the rationale. So that should not come as any surprise.

The implementation guidance is that the application submission system should be [all feasible] algorithmic checking of TLDs, including against the root zone label generation rules and ASCII string requirements to be better ensure that only valid ASCII and IDN TLDs can be submitted. A proposed TLD might be algorithmically found to be valid, algorithmically found to be invalid, or verifying its validity may not be possible using algorithmic checking. Only in the latter case, when a proposed TLD doesn’t fit all the conditions for automatic checking, a manual review should occur to validate or invalidate the TLD. In other words, when someone is submitting their applications, they should, to the best of the ability of ICANN, be on notice as to whether they could apply for that string or not.

Rubens makes the point that all emoji submissions should be answered with a sad face, which I guess is better than some sort of other anger emoji.

Justine puts—yeah—“It’s SAC 095 that is the paper that talks about not allowing emojis.”

Okay. That takes us through the security and stability. If you want to read the rationale and the background SSAC papers on this, you should do so. I can’t remember if we have it in here, Steve, but I think we should also include the ... Well, we do include the references to the papers and to the charts, so I think that’s good. We do point out the dependencies here which are involved in ... We talked about registry system testing, IDNs, applicant reviews, financial registries systems, and name collisions as it related to DNS stability.

We will talk about name collisions on the next call. So, for the next call, please do read not only the section that we will circulate shortly but also, if you have time, the paper that was submitted that I sent out to the group from the NCAP discussion group. Again, it is just in draft form, so we cannot take that as the final be-all/end-all. It will be discussed this week, I’m sure. I think the NCAP discussion group is meeting on Wednesday. Regardless, it will be the subject of the NCAP discussion group, so that document may change. I will say that, on the mailing list, there were some that disagreed with the recommendation of the independent consultant. The independent consultant recommended or basically stated that they did not necessarily see a justification for doing Studies 2 and 3. It is an interesting recommendation, but it is not one that is supported by everyone in the NCAP discussion group. So it does make for interesting reading.

With that said, if we can post the next time.

Jim states, "Are we discussing the NCAP report or our work?"

Jim, we're going to talk about our work, but our work does reference to the NCAP work. So, for some background, it is good to read what the NCAP has been up to. So some of those may relate to the work. As an example, one of the recommendations we were talking about is whether the delegations should wait on the NCAP doing Studies 2 and 3, but it is not necessarily a given, if you read the independent consultant's report, that there will be Studies 2 and 3. So I think there are certainly reasons why we should at least be aware of what the NCAP group is working on: because they will have an impact on our recommendations. I hope that makes sense, Jim. So it's not that we're substituting our judgement for the NCAP group. It's just an awareness of what's going on there.

The next call will be—oh, sorry. Alan, go ahead.

ALAN GREENBERG: I was just going to close up the discussion on delegation rate. With a quick look at contracting dates, which is close to delegation rates, it looks like we hit about 500 a year. So 1,000 was not that far, and [as] we get better, we might hit 1,000 in the future. Just a point of ...

JEFF NEUMAN: Thanks, Alan. Remember, too, that things may speed up a little bit because, if we do have the RSP pre-evaluation program in place, the technical evaluation for a number of strings may not have to

take place during the normal application window. So that may also speed up.

ALAN GREENBERG: Yeah. Just noting that it was a factor of 2 out from 1,000, not a factor of 10.

JEFF NEUMAN: Yeah. Thanks. That's a good point.

And Maxim states, "We don't know how many years the total NCAP effort may take."

That's true, Maxim, but let's save that for the discussion on Thursday, where we will also have the Co-Chairs from the NCAP discussion group there. I will ask them to provide a brief update on the work at the beginning of the call so that we can consider that update along with our work that we've done on the subject.

With that said, thanks, everyone. I know we made a lot of progress here. We covered a lot of technical materials here. I appreciate everyone for being on and for making progress on these issues. Thanks, everyone.

[END OF TRANSCRIPTION]