

---

**ICANN Transcription  
IDNs EPDP**

**Thursday, 20 January 2022 at 13:30 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 of the call are posted on agenda wiki page:

<https://community.icann.org/x/m4H3Cg>

The recordings and transcriptions of the calls are posted on the GNSO Master Calendar page

<https://gns0.icann.org/en/group-activities/calendar>

DEVAN REED:

Good morning, good afternoon, and good evening. Welcome to the IDNs EPDP Call taking place on Thursday the 20th of January 2022 at 13:30 UTC. 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 telephone, could you please let yourselves be known now?

All members and participants will be promoted the panelists for today's call. Members and participants, when using chat, please select Everyone in order for everyone to see the chat. Observers will remain as an attendee and will have View Only to access.

Statements of Interest must be kept up to date. If anyone has any updates, please raise your hand or speak up now. If you need assistance updating your Statements of Interest, please e-mail the GNSO secretariat.

All documentation and information can be found on the IDNs EPDP Wiki space. Recordings will be posted on the public Wiki space shortly after the end of the call.

---

*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.*

---

Please remember to state your name before speaking, for the transcript. As a reminder, those who take part in the ICANN multistakeholder process are to comply with the expected standards of behavior.

Thank you, and over to our chair, Donna Austin. Please begin.

DONNA AUSTIN:

Thanks, Devan, and welcome everybody to this week's IDN EPDP call. Just in case you're wondering what you wandered into, we were just having a little bit of chitchat with Justine and about how she got started in ICANN. But we're not expecting everybody else to give an overview during this call of how they got involved. But I realized that one of the challenges with not being face to face is we don't get the opportunity for just a little bit of chit chat to find out a bit more about people. So I kind of like to open the room early and see who comes in and just have a bit of a chat. So that's what we were doing.

So with that, substantively, we'll be talking the question A5 and A6 today. And if we get time, we'll move on to A7, A9 and A10. I don't think that will be the case. I think it's a little bit optimistic. I did mention a couple of weeks ago that we'd be looking into whether reordering our charter questions makes sense or not. We're still reviewing that. But the idea is that we would only suggest changes to the order in the event that we think it's substantially helpful. So we're still reviewing that. And hopefully we'll have some feedback on how we want to approach that in the next week or two. But for now, I think we're agreed that we will work our way through the topic A questions and draw a line under those before we would

---

make any changes to the order. So that's the status of what we're doing with that.

And Farrell, I know you're on the call. Are you okay to provide an update on the letter that the Board wrote to Council about six weeks ago about the IDN guidelines?

FARRELL FOLLY:

Hello. And as Ariel has just said, I will just give you a brief overview about the discussion that happened between the GNSO Council and the Board regarding the deferral of the implementation of IDN guideline version 4.

So basically, Council has agreed with the Board suggestion to continue to defer the item that overlap. The work of the IDN PDP should continue. So the Council responded to the Board and identified some items in the guideline, mainly the points six 6, 8, 11, 12, 13 and 18 and working on that. And actually, the Board didn't identify those points but asked the Council to do so. And that has been done. And the Council now is expected to send the letter by the end of this week. Basically, after we have the GNSO Council meeting tonight.

So this is just an informational news and this is not an action item per se for this EPDP, but that means Anil who are respectively the liaison or who can transfer this information to the ccPDP 4 and the CCNSO can now have this information about what is going on within the GNSO Council regarding this issue and a discussion with the Board. So thank you very much. This is what I have to

---

say regarding this issue. I hope I didn't forget any critical information.

DONNA AUSTIN:

Thanks for that update, Farrell. So does anybody have any questions regarding that, just noting that it will be an item for discussion on the Council call in a few hours from now? Okay, I don't see any hand. So I think we can move past that. And Farrell, if I could ask you that if there's any update following the Council meeting, if you could just provide that to the list to keep people informed. That would be appreciated. And it may also be something that our liaisons to the ccNSO work, if you can just, keep them informed of what's going on in that regard as well. I'm not sure how closely the CC group is following it, but it probably has potential impact on their work as well. I'm not 100% sure on that.

Okay, so we want to start with charter question A5, which was what we had a lot of substantive discussion with the SSAC about last week around whether there should be a ceiling mechanism or not related to variants. But before we get into—and I'd like to have a discussion about how folks feel about the discussion that we had with SSAC and whether it's helped us in our thinking on A5.

But before we do that, I believe Sarmad is going to provide us with some information that will, hopefully, help us as well with charter question A5. There's a little bit of data that Sarmad is going to provide us with on scripts, how they treat variants with the individual scripts. So Sarmad, if you're okay to present, I'll hand it

---

over to you, and then we'll come back to the discussion on A5.  
Thanks.

SARMAD HUSSAIN: Thank you, Donna. I think Ariel was going to start and then I'll add on. Or are we jumping directly there?

DONNA AUSTIN: Now, sorry, I may have misspoke. So Ariel, if you're going to open this and then hand over the Sarmad, then please go ahead.

ARIEL LIANG: Thanks, Donna, thanks, Sarmad. I just have a very brief introduction of what Sarmad is going to present to the team. So in order to understand the charter question of whether there should be a ceiling value or other mechanism to ensure the number of delegated TLD variant labels remains small. Sarmad did a data collection exercise to check all the scripts in the current RZ LGR as far as the scripts that will be incorporated in the next version of RZ LGR, to understand whether there's already existing mechanism to reduce the number of allocatable variants. So in that way, we may have some kind of numbers or baseline to take a look at and see whether the information he collected will help in any way answer this charter question. So Sarmad is going to present you the findings. And hopefully, this will be helpful as another point to address this charter question, in addition to what the team discussed with the SSAC. So that's a very quick introduction and I'll turn over to Sarmad for the presentation.

---

SARMAD HUSSAIN:

Thank you, Ariel. And thank you, Donna. Basically, what we were asked to do was to get some data on how many allocatable variant labels are generated by the different scripts, which are being integrated into the root zone LGR. So you can see that this lists all the active generation panels, most of them have finished, and some of them actually are finishing with basically the proposals either very recently going through the public comment, or currently Myanmar is the only one left which is open for public comment.

So if you look at these lists, there are four scripts, Georgian, Gujarati, Lao and Thai which don't have any variant labels at all, which obviously don't pose any problem or challenge in this context. We also have a longer list in the middle, not going to read all the names, but these are scripts which have variant labels, but none of those variant labels are allocatable, they have actually defined all the variant labels as blocked.

So that also does not create any challenges, because what we are focused on is looking at the challenge of, I guess, seeing how many allocatable variants are produced, and then possibly, whether they can be reduced or managed.

So the scripts in this context which are relevant are the ones in the final column, which are Arabic, Bengali, Chinese, Greek, Latin, Myanmar and Tamil. So let's look at them in a little more detail. Sorry, let me stop here. I see a hand by Anil. Anil, please go ahead.

ANIL JAIN: Thank you so much. I just have a question here that this list into the three format, does it only cover gTLDs, or does it also cover ccTLDs as far as the variants are concerned? Thank you.

SARMAD HUSSAIN: Same would apply to both ccTLDs and gTLDs. So this is basically the list covering the root zone LGR. And it's eventually going to be based on the policy for the gTLDs and ccTLDs on how that is incorporated. So eventually, the final decision will be done by the policy. This is just the contents of the root zone LGR. So there's a question in the chat by Donna. Please go ahead.

DONNA AUSTIN: Yeah, thanks, Sarmad. I wasn't sure whether you'd see the chat or not. I just wonder if you can give a little bit of context as to why we have this situation where some have no variant labels, some have no allocatable labels, and then others do have allocatable variant labels. Is there a technical reason for that? Or is it a kind of a policy decision that's made within the generation panels themselves?

SARMAD HUSSAIN: There are two distinctions here. One between scripts which have variant labels versus scripts which do not have variant labels, so that's column one versus column two and three. So that is normally dependent on whether the community thinks that any codepoints in their repertoire should be considered same. In some

---

---

scripts, those script communities have found two different code points in Unicode which the community thinks that are the same. So therefore, they define it as variants.

In some other scripts, like Georgian or Gujarati, the script community thought that all the codepoints which are in their repertoire are quite distinct, and they're actually not the same. And therefore there are no variants within that script. Or there is no matching within between that script and any other script as well. So for example, many of you are familiar that we have matching between Cyrillic script and Latin script. So there are also those cross scripts similarities, sometimes, identical glyphs. But in these scripts, they have looked at the script internally as well as cross script, and they found no two code points are the same so they actually concluded that there are no variants. But that's obviously, as you can see, the smaller set.

Then we come to the question that for those scripts which have variants, whether those variants are allocatable or not. So the primary reason why a variant would be made allocatable is not the security argument, but a usability argument. So the security argument is already satisfied when you say that you have variants, because once you have variants, and if the policy says at least what's being recommended is that variants go to the same applicant, then the security aspect is contained.

However, the usability aspect is not contained. The usability argument comes that if there are certain variants, and you are in one geography, or territory or country which is using one variant, and you are in another geography, territory, or country which is

---

using another variant, then if you block all the variants, all the different communities cannot access that TLD anymore.

So, we are very familiar with simplified Chinese being used in China. And then you have traditional Chinese used in, for example, Hong Kong, Taiwan, some other places and if you really want Chinese to be used across the board by everyone, so there's a usability argument, you will need to make the variants allocatable. But if you're, for example, talking about Cyrillic or Armenian, where they think that there are no variants within the script, but for Cyrillic, there are some matches across the script with Latin, for example, or Armenian, those cross script variants are not allocatable because there is no usability argument there.

So in a way, you can label these three columns as the first column, which says no variant labels, that there is no security issues and therefore no variants are required for the purpose. This middle column means that there are some security potential security issues because some code points can be considered same. But there are no usability issues or reasons for the middle column. And then the final column, you have security challenges, which can be caused by "same" codepoints. But some of those same codepoints also have some usability requirements across different populations. So that would be one way of how one could distinguish between the three columns. I hope that answers your question.

DONNA AUSTIN: Yeah, that's really helpful. And I see that Hadia has her hand up. Hadia.

HADIA ELMINIAWI: Thank you so much, Sarmad. Thank you, Donna. Thank you for this question, because actually, it clarifies a lot of things. And I would ask, Sarmad, if you could please repeat again the security aspects. So as I understood from you, the security problem with the third sector is that you have similar codepoints that could potentially be allocated. But then I did not get how do we overcome this.

SARMAD HUSSAIN: So the security challenge is the same for column two and column three. There's no difference. And the security challenge arises from the fact that there could be two different codepoints, either within the script or across different scripts, which are not similar, but actually same, quote unquote, same which the script community considers as indistinguishable from other codepoints. So by declaring such codepoints as variant codepoints, the security angle is addressed for both column two and column three. The only difference in column three—

HADIA ELMINIAWI: One and two you mean.

SARMAD HUSSAIN: No, two and three. So column one would not have any security challenge at all, because those—

---

HADIA ELMINIAWI: Yes, exactly, and two also has no, problem, but three—

SARMAD HUSSAIN: Two has problems because they defined variants, but they defined all the variants as blocked variants.

HADIA ELMINIAWI: But there is no problem because it's not allocatable.

SARMAD HUSSAIN: Right, by defining variants as blocked, the security challenge is addressed. That's what I'm saying.

HADIA ELMINIAWI: Yes, yes, exactly. So I understand this—

SARMAD HUSSAIN: To make them allocatable from blocked, that motivation does not come from the security challenge but from the usability need. So that's the unique thing in the last column, that even though by defining variants, the security challenge is addressed, the usability challenge is not addressed and that's why for these scripts, they want to make some of their blocked variants allocatable.

HADIA ELMINIAWI: Sarmad, again, I understood what you said, that no variant labels, there is no security issue here. No allocatable variant labels, we don't have a security issue, because the security issue is already

---

addressed. Having allocatable variant labels, the security issue still remains.

SARMAD HUSSAIN: No, there's no security issue anymore. But the reason we want to have them as allocatable variants is based on their usability requirement. Security is already addressed by having blocked variants, or having even allocatable variants. If you have variants allocatable or blocked, the security challenge is addressed.

HADIA ELMINIAWI: Okay, so if we say those are allocatable, that means that we have no security issues with it, right?

SARMAD HUSSAIN: Yeah. Whether blocked or allocatable, doesn't matter. The security challenge gets addressed through that variant label process. Quoc, please.

QUOC PHAM: Hi. I think this is a really interesting topic. I think the way we should look at it is variant evaluation at the root level should be there to protect against similar looking TLDs. I think that's the motivation behind all of this, which is the whole point of it.

I think the distinction between none non-allocatable and allocatable adds a little bit of confusion, but the way I see it in terms of unpacking it is like this. Using Chinese as an example, I'll post it inside chat right now, those two characters 39DF and

---

64D3, they're listed as variant characters, meaning that the presence of one eliminates the presence of the other because both of them aren't unique. They're equivalent.

So if you look at the label towards the left-hand side of what I've pasted, that's 39DF three times in a row and 64D3 is the next IDN label pasted three times in a row. So in terms of evaluation, if I was to apply for the TLD to the left, then the TLD towards the right cannot exist independently. I think the idea of allocatable or not in this context of root zone evaluation is, should label two be granted to the registry operator that applied for label one because there is an allocatable variant? That's a question that needs to be answered.

With regards to the variant question altogether, though, from a computational point of view, when a TLDs is applied for, and it goes through the root zone evaluation of the TLD that's been applied for, is it unique in a sense that there's nothing else that exists in the root? Where canonically, meaning that if unpack each single character in the string of the TRL, is that unique? Right? Does it stand on its own? And if it does, if it is unique, then you get in the next phase, you say, okay, what do we do with it? But if it's not unique, meaning that there's something that already exists, then it should fail in terms of the application itself.

I'll paste another example in here for context. Cyrillic small letter C, Cyrillic small letter O, Cyrillic small letter EM. That looks like com. So if I was to apply for an IDN in Cyrillic in that particular format, that should be rejected because .com already exists, right? That should go against the Latin evaluation of—what was it, cross script? I think somebody said cross script variants.

SARMAD HUSSAIN: Yes.

QUOC PHAM: So that is the scope of the root zone evaluation for IDN TLDs. But [even if you] throw out the context of IDNs, even if you apply for an ASCII TLD, it needs to go through an evaluation to see if that ASCII TLD being applied for is unique in the sense of the character mapping of this entire library that's been presented in front of us.

DONNA AUSTIN: Quoc, sorry to interrupt, I think you might be jumping ahead of a discussion we'll have a bit later. So is it possible to hold your thoughts and we'll let Sarmad get through the rest of this and then we'll come back to that discussion when we get into the discussion proper about the charter question?

QUOC PHAM: Of course, sorry about that.

DONNA AUSTIN: Yeah, no worries. Thanks. Sorry, Sarmad. So if we can just let Sarmad go through and finish this presentation, which is to provide context that will help us answer the questions. And we'll get to some of the stuff that Quoc was talking about. Sorry, Sarmad. Go ahead.

---

SARMAD HUSSAIN: Thank you, Donna. So our relevant scripts are those which are in the rightmost column. Those are the ones which have allocatable variants, and the question to see is whether there is an upper ceiling or upward bound on these variants which are generated, and whether they match the possible expectation of too many allocatable variants or not. Next slide, please.

So we looked at each of the scripts separately, and we wanted to test them against some data. And obviously, one can generate any set of labels to test these scripts. What we did was that we actually went through the existing list of top-level domains which are delegated, gTLDs and ccTLDs and we ran them through the root zone LGR to see how many variant labels are created. And based on that, we're making some comments, this is just data for you to review and then of course, discuss in the context of the larger question.

So for Arabic, one thing to note is that there is no upper bound on the number of variants which can be created. So it will depend on the codepoints you're selecting and the length of the label, and it could actually potentially create a large number of allocatable variants. So we are only discussing allocatable variants here, right, we're not discussing blocked variants. Blocked variants can be very numerous as well, but that's not relevant, because they will actually be blocked and do not come in the discussion.

So for Arabic, there is no general upper bound. It is a function of the codepoints included as well as the length of the label. However, when we ran the TLDs in Arabic which are currently delegated, we found that the allocatable variants ranged from

---

anywhere between one variant label to nine variant labels for these delegated top-level domains.

For Bengali, or Bangla, a variant letter for one of the languages is allocatable, while the others are blocked, so it's just one character which has an allocatable variant. And there is actually a rule which says that if that character occurs more than once in a label, either you can have only one type of that character or the other type of that character, but you can't mix those characters, very similar to the Chinese example which Quoc shared in the chat, where the two characters for example, which are variants, so you can formulate a label with one kind of character or codepoint getting repeated or the second one, but if you don't mix them, then there are only two options available. So there is an upper bound of two for any of those for Bengali.

Chinese, when we ran the actual labels through which are delegated in the root zone, the variant labels we got was from one to three variant labels. So it's also a reasonably contained set. Generally, the way the LGR is designed is that the LGR will allow for the applied for label and then all the possible simplified versions, all simplified version of that label, and then all traditional version of that label. And generally that is bound to three labels. But in some cases, where there are variants, for example, within all simplified or all traditional, there can be a few more variant labels as well.

Greek, there are at most four labels which can be allocated. There are two allocatable kind of codepoints or sets of codepoints in Greek. They have some vowel marks which have [inaudible], so that accented versus unaccented vowels. So the Greek

---

generation panel considers them as variants of each other. And then in Greek, we have a final and nonfinal sigma. So they actually have a rule which limits the number of variants which can be allocatable. They say that you can apply any arbitrary mixture of these vowels and unaccented vowels and sigma and final sigma. But then the variants which are created can only have either all unaccented vowels or only the final sigma form. So based on that, they can be at most four allocatable variants in Greek. Next slide please.

For Latin, there are just two sets of variants, one which is created between sharp S and double S for German, and one with dotless and I. And Latin GP also creates basically a rule which does not allow arbitrary mixing of these. So one could get the original label which one applies for and then one which has all sharp S'es converted to double S'es, or all dotless I's converted to dotted I's. So that also creates a maximum of four labels.

Myanmar also has a similar set where—and this is based on the current public comment version. So this is not the final version yet. But what they've done is divided codepoints in two sets and there are variants across two sets, but they actually have a rule that the codepoints can come from one set or the other set, and the original label. So, maximum of three allocatable variants which are allowed.

Tamil also has a similar context. So for Tamil, in summary also, there are two possible variant labels allowed. So in summary, if you go through all these, most of these scripts have been able to, by the nature of these scripts and the by nature of variants, they have put a reasonable ceiling on the number of variants. But if you

---

go one slide back, there are at least two scripts, Arabic and Chinese, which may not actually have an upper limit and so may create an unanticipated number of variant labels. Let me stop here and see if you have any questions. Otherwise, it's back to you, Donna. Thank you.

DONNA AUSTIN: I see a couple of hands. Satish.

SATISH BABU: Thanks, Donna. Thanks, Sarmad, for the presentation. It's exceedingly useful. I just have a question relating to the source of this work, especially the first classification into three buckets. Is there an official document somewhere? Has it been done by the IP or by your team? Can you clarify?

SARMAD HUSSAIN: We just went through all the root zone LGR proposals by all the script communities.

SATISH BABU: Okay, so this has not been published. This is new work.

SARMAD HUSSAIN: It is in a way published individually in the proposals, but we've just collated that information for GNSO IDN EPDP.

---

SATISH BABU: Great, thanks very much.

DONNA AUSTIN: Okay, so we have Justine, and then Hadia.

JUSTINE CHEW: Thanks, Donna. Sarmad, could you perhaps clarify for me—Well, I see this on slide four now, where the proposals for Japanese, Latin and Myanmar, are still being finalized. So I wanted to ask specifically about the Latin script proposal. Has what you presented taken into account the public comments that have been received, or not? Thanks.

SARMAD HUSSAIN: I see Michael's hand. Michael's part of the Latin GP. It may be good to have him come in on this if that's okay, Donna.

DONNA AUSTIN: Absolutely. Go ahead, Michael.

MICHAEL BAULAND: As Sarmad said, Dennis and I were part of the Latin GP. And to your question, yes, we have taken those public comments into account and they do not change anything relating to these allocatable variant definitions here. So this will definitely not change after we apply all public comments. Thanks.

---

JUSTINE CHEW: Very good. Thank you, Michael.

DONNA AUSTIN: Okay, thanks, Justine and Michael. Back to you, Hadia.

HADIA ELMINIAWI: Thank you. And if we could go back to the first slide. My question to you, Sarmad, so the label generation rules does not [permit] any variant mapping that may pose security concerns. And security concerns are addressed by the security panel. Thus, all allocatable variant labels are deemed allocatable because the security panel said that they are safe. And my question here, is there a reason to think that security issues might change as we add more scripts or for any other reason? Is there a reason for us to think that security panels will need to assess those allocatable variants for any reason in the future?

SARMAD HUSSAIN: So let me, I think, qualify my statement. Or maybe I should probably, in a way, correct it. When I say that the security part is addressed by creating the variant labels, that's in a limited, I guess context, not all security questions are addressed. They're addressed from just the perspective of how the scope which is defined by the root zone LGR procedure, which actually asks the generation panels to create variant codepoints which they consider are interchangeable, the codepoints which are interchangeable, it asked them to define them as variants. And it actually lays a very strict definition there, it says that only the obvious cases should be identified as variant labels.

And what the root zone LGR is doing is only capturing that part. There is much more to security. And I'm sorry if I gave the impression that everything is addressed through this. That's not the case. Everything which is related to same labels is addressed through the root zone LGR process. But there is obviously more to security, which relates to obviously other dimensions, even some of the string similarity issues which are not obviously, quote unquote, same as per the root zone LGR definition may also sometimes have some implication.

So the root zone LGR only solves a very limited, defined problem of, quote unquote, same codepoints. Anything related to security which is caused by other factors, for example, similarity issues or some other stability or security issues, those have to be dealt by the DNS, obviously, other panels like DNS stability panel, or string similarity panel. Thank you.

DONNA AUSTIN:

Okay, thanks, Sarmad. I don't see any more hands. So thank you for the presentation. I think it was really important because it puts some context to—we know the bounds of the discussion that we need to have around whether there should be a ceiling or not as it relates to charter question A5. So I think this is really helpful to understand that what we're really dealing with here is a small number of scripts. And within that small number, the seven, they already have identify their ceiling within the LGR. So I think we're kind of left with two. So I think that was really helpful context for the discussion.

---

So, Ariel, could you just put the A5 question back on the screen, please? Okay, thank you. Well, it may or may not, Jeff. We need to have that discussion. So I want to open it up now. And maybe it would have been helpful to have this context before we had the discussion with the SSAC, but there you go. I just want to get folks' thoughts now on how they're thinking about the charter question A5 in the context of what we just heard from Sarmad and also the discussion that we had with SSAC last week. Jeff, go ahead.

JEFFREY NEUMAN:

Yeah, thanks. To me, it appeared that—at least I was under the impression that, at least with the SSAC members that were there, that we were of the same kind of thinking, which was not that we come up with some kind of arbitrary number as some sort of artificial limit, but rather focus on questions that should be asked to the registry and evaluation criteria that should be used by ICANN or whoever the evaluators end up being, in order to understand that the registry can, quote, manage—and I'm using to manage because I think that was a term that was used by the SSAC—can manage the different possibilities and different issues that arise when you are delegating or when you do have different variants that are delegated, which includes not just how the registry itself is going to manage its own processes, but how it's going to communicate with the registrars and ultimately with the registrants about how they can manage the number of names that they may be forced to own or deal with, or whatever it is.

So again, I really haven't kind of changed the view based on any of this information that we should not be, as a policy matter,

---

essentially protecting entities from themselves in the sense that registries are going to apply for what they think they can handle. And then they need to be evaluated as to whether ICANN or the evaluators agree that they can handle it. And if they can, then who are we to say they can't have more than one variant? Thanks.

DONNA AUSTIN:

Thanks, Jeff. So I'm really interested to hear how others feel about the way that Jeff has characterized what he thinks could be a solution. And also interested to hear from—because I think this was something that the SSAC brought up, is that there's no technical solution to the bundling, so that is something that needs to be sorted out. So I'm interested to hear from maybe the registry and registrar folks that we have on this call about whether they have any insight into some of the challenges that might arise in that bundling exercise. And I see that within the chat, Michael fully agrees with Jeff. Quoc, go ahead.

QUOC PHAM:

Hi. First of all, sorry about my little off topic rant before. I think, in my opinion, when you apply for A TLD and there's an evaluation on the TLD, you don't get automatically granted, I think, at all, any variant TLD based on the evaluation.

And the reason why I feel that way is because the applicant looks at the TLD that they want to apply for. And it's a purposeful combination of characters that's chosen for whatever reason. And it just so happens to be that there is a computational evaluation on

---

it to say, “Well, okay, we're going to protect that for a security sense where a similar looking TLD cannot be applied for.”

Now, if ICANN wants to offer the applicant to say, “Hey, listen, to do you want the variant TLD of what you've applied for?” From a registry perspective—now putting on a registry hat here—number one, the solution to it is to be decided, but from a DNS perspective, that's a separate design. So it's another TLD that has to be managed, that has to be deployed on the DNS network that the registry has chosen. And there's maintenance of two zone files for two TLDs.

So it's just another TLD on the registry at the lowest level. There's a bunch of business rules that sit on top of that in terms of solution, in terms of how you manage registrations. The term “bundling” has been thrown around. For example, if you register a domain name in TLD A, suddenly, the domain name in TLD B is also present and it resolves to the same name servers or whatever the case is.

But it's still two separate TLDs basically at a DNS level that the registry has to facilitate and operate. And the cost, obviously, doubles or triples depending on how many variant TLDs exist with the application.

And then also on top of that, then we have to look at the second level and also the problem at the registrar. There's obviously another problem there to solve. So in summary, I don't think an application should automatically be granted a variant TLD. I think an application should be protected from having other TLDs that look similar based on variant rules to be prevented from entering

---

the root zone. But granting an extra TLD, variant TLD automatically, I think it's not necessary.

DONNA AUSTIN:

Thanks, Quoc. And just a question that comes to mind. And maybe others can think about this. But is the intent of a variant label, if it is allocated, to have the same user experience as if you went to, I think, Jeff, you called it the primary TLD last week? So what's the purpose of the variant and whether that will be used or not? Edmon, go ahead.

EDMON CHUNG:

Yeah, just put up my hand to add a little bit what Quoc said or moderate a little bit. I think in many cases, that's probably correct. But in general, the concept of IDN variants is not necessarily lookalike. And I think you put out some examples from Chinese that is clearly not so much lookalike.

But in Chinese, I want to remind everyone that—and sorry, I misspoke on this particular statistic before, but I got the right number now, up to 5% of users would type in the variant itself. And right now they go nowhere for the gTLDs that are in place.

So, one out of 20 users on the Internet that is typing in Chinese domain names would go to the variant. And that statistic's taken from CNNIC and TWNIC. So I think it's not just for protection. It's also for the usability of these IDNs, at least in the Chinese case.

---

DONNA AUSTIN: Thanks, Edmon. Jeff.

JEFFREY NEUMAN: Yeah, so I think we kind of got off the question. And I agree with Quoc in the sense that of course, we shouldn't be automatically giving registries a variant. I don't think that's the question. Question is whether there should be a limit on the number of variants that a registry can have or can, quote, manage. And I don't think—again, I come back to even as Quoc said, it has to be separately evaluated as to whether the registry can manage it, it shouldn't be automatic. But by the same token, we shouldn't just impose any kind of limit.

DONNA AUSTIN: Thanks, Jeff. Sarmad.

SARMAD HUSSAIN: Thank you, Donna. I just wanted to add one more detail, more of a factual detail which I think may also be relevant to this discussion. So when we are generating allocatable variants through the root zone LGR process, even though the generation panels have worked hard to put rules to contain labels or allocatable labels which are, in a way, to some extent, extraneous, the process still, in some cases, over generates.

When I say over generates, what I'm saying is that when people are using a label, they are only looking at it from the language perspective. So for example somebody who's using Arabic script, for example, may be looking at the label from an Arabic language

---

perspective, or Urdu language perspective of or Persian language perspective, they're looking at labels from their perspective.

But since this root zone LGR is done at the script level, sometimes it generates variant labels which are mixing characters across different languages. And those are characters which, for example, no one language community would be using because it has some characters from their codepoints, from their set of letters, but they have some extra letters which is used being used by some other language community.

So there is some level of over generation which is not directly motivated by the usability but that's sort of a byproduct of making some of these solutions work. So I just wanted to share that information as well, that not all allocatable variants are equally usable. Thank you.

DONNA AUSTIN:

Okay, thanks, Sarmad. And I see that Justine has put in the chat that even if we don't think a ceiling should apply, we maybe should consider concretizing the criteria by which registries apply or requests are evaluated. And I think that's consistent with what Jeff has suggested. And there seems to be some agreement to that in chat. And I would also say that that probably addresses the second part of this question about, should additional security and stability guidelines be developed to make variant domains manageable at the registry, registrar and registrant level? And I think what I'm hearing is that the answer to the second part is yes. And I guess the question for us is, do we want to try to put some

---

suggestions around that, whether it's a policy recommendation or whether it's implementation guidance? So Jeff, go ahead.

JEFFREY NEUMAN:

I just want to provide a clarification. I agree that additional guidelines should be developed. But I'm not 100% certain we should be calling them security and stability guidelines, because the SSAC said that there's no kind of technical rules. I know it's a little bit nitpicky. But I think my answer to that would be additional guidelines should be developed, but more from a user perspective, as opposed to from a security and stability perspective.

DONNA AUSTIN:

Okay, so we can bracket the security and stability, but I think there's agreement that we definitely should provide you—whether it's a recommendation or implementation guidance on guidelines to make variant domains manageable at the registry, registrar and registrant level.

So I think what I'm hearing is—and I think it was really helpful to get Sarmad's presentation to understand that what we're talking about here isn't an enormous—it's not every script is going to have this problem, because most of the scripts have solved the problem that we're trying to solve here already. It only applies to seven scripts, and within that seven, I think there's only two that feed into something that might be considered a larger number. So that was really helpful.

---

So I think what we're hearing here, what I'm hearing, and I think there's support for, is that we're not recommending a ceiling value, and part of our rationale will be the information that Sarmad has provided to us and also, we think that there is some support for the—this is something for the market to provide as well.

So, unless anyone else has any other contribution on this topic, I feel like we're in a pretty good place on this. It may be that we'll have to come back and put some meat around the guidelines or at least some ideas in place for the guidelines, and maybe we can look to the registry and registrar members to help with that. But otherwise, I think we're in pretty good shape. Any last words from anyone? okay, I don't see any, so this is great. Thanks, everybody, for the contribution, and we can now move on to A6.

Okay, so it did take us a little while in the leadership team to get the draft language together on A6, so thank you for your patience. I hope everybody's had an opportunity to read through the text as we proposed it. And what we want to do here is get some feedback. So whether there's agreement on the language or not, and then we can discuss the why or not, and then we had a number of questions that followed the language that we want to discuss as well. So I just want to open it up. I think Sarmad had a couple of questions in the text directly. So can we go to that, please, Ariel?

Sarmad, did you want to just provide some context to your question here? And then we can have some discussion about how folks feel about that.

---

SARMAD HUSSAIN: Sure, thank you, Donna. So when we're talking about variant labels for top-level domains, they will fall into three categories once all this is operational. There will be a set of variant labels which will be blocked and stored away, and nobody would be able to use them. There will be also a smaller set which will be allocatable potentially to the same applicant. And they can then eventually decide if they want to get them delegated or not. And then there will be possibly a smaller set which will be delegated variant labels.

So, obviously, when a variant label or string is delegated in the root zone, it has a very high stability bar, meaning that it should remain delegated, and it should not—unless it's a very extreme reason, it should not be undelegated or changed from being a variant to a non-variant.

Allocatable variants are, again, those which are in reserve with possibility of getting delegated, but they're not actually in operation right now. And therefore, changes in them may not be as critical as changes in the delegated variant labels.

And then if you go to the third category, which is the blocked variants, since they are stored away and not for use by anybody, their change in, I guess, disposition or change in that part of the set is even less critical, in a sense, because they were blocked away and they were actually not even possible to be allocated or delegated.

So I guess my question was that, in recommendation 1.4 and recommendation 1.5, we talked about variants. And I guess the question was, do we want to differentiate these three levels of

---

variants—variants which are delegated versus allocated versus variants which are blocked—and see whether one or all three of those subcategories are applicable to 1.4 and separately to 1.5? And so that's sort of the question I was raising. Thank you.

DONNA AUSTIN: Thanks, Sarmad. Justine and then Michael.

JUSTINE CHEW: Thanks, Donna. Sarmad, if I could ask you to—and maybe I didn't grasp in total what you were saying earlier. But if I could ask you to explain the implications of the different buckets that you refer to in how it applies to recommendation 1.4 and 1.5, please.

SARMAD HUSSAIN: Sure. So suppose if variant TLD was delegated, and it is in the root, then if for some reason, there is a change in root zone LGR which says that—and the change can be of multiple types. So let's assume that the change is saying that that particular variant should be blocked, or another change could be that that variant is for some reason no longer a variant. Then we would want to grandfather it, because it's already in the root zone since it was delegated.

But if it is a variant which is blocked, which means that it's not being used by anybody and will not be used by anybody in the current state, and if it goes around and maybe becomes allocatable, that doesn't really cause, for example, a significant quote unquote issue, or at least as significant an issue as trying to

---

change something which is in the root zone and delegated. So the category of variants has some implication on how a change in that category can have—how much impact that change can have on the stability of the root zone. Thank you.

DONNA AUSTIN: Does that answer your question, Justine? I'll take that as a yes. Michael, go ahead.

MICHAEL BAULAND: To answer Sarmad's question, or at least to voice my opinion of a possible answer, is that I would suggest restricting this to allocated variants but ignoring the blocked variants. As Sarmad already said, the blocked variants are just a technical thing and registry will never be able to use that. So there's no harm done if that got taken away by root zone label generation change. But an allocated variant, even though it's not yet delegated, I think the registry already paid for it, applied for it, then it got assigned to the registry and they may have plans to just delegate it a little bit later. And for that reason, I'd say that both delegated and allocated variants should be considered here. Thanks.

DONNA AUSTIN: Thanks, Michael. Thoughts from other folks on this one? Anyone agree with Michael's suggestion? Dennis agrees with Michael. And Satish as well. Okay, so I think maybe Ariel's ahead of us here. I think what I'm hearing is delegated and allocated. So we need to be specific about what we're talking about here. To Sarmad's point, I think taking out any ambiguity would be helpful.

---

So for delegated and allocated, the intent would be for backward compatibility. But with blocked, not so important. I think that's where we are. Okay, great. So thanks for raising that, Sarmad, and for the discussion. Sarmad, go ahead.

SARMAD HUSSAIN: Just one small edit there. I was going to suggest allocatable, but I think Ariel already made that change. Thank you.

DONNA AUSTIN: Okay, great. Thank you. So can we have some feedback on the text as drafted? Is there anything that people thought we've completely missed the mark on this, or are folks comfortable with how we've captured the discussion and the recommendations? Justine's asking if we can use variant labels instead of variant.

Okay, so we will hold this text open for another week. I think the intent was we'll always have a two-week window for folks to consider text. But if there's nothing that comes back, then we'll assume that we're okay, notwithstanding that we need to get to those other questions that are yet to come. Sarmad.

SARMAD HUSSAIN: Yes. I also had, I guess, wanted to ask the recommendation 1.6, are we discussing that now as well?

DONNA AUSTIN: Yeah.

SARMAD HUSSAIN: Okay, so the recommendation 1.6 actually suggests that we go and potentially update the root zone LGR procedure itself. So in the context of not being able to retain backward compatibility, we need to make some changes in it to define the circumstances where this could be possible.

I just wanted to make two observations in this context. First of all, that the root zone LGR procedure already has a built-in stability clause. It works as a principle, which actually does build that in already, to a large extent. I guess, not to a very detailed level of specification, but as a guiding principle. And then eventually that principle is interpreted by the generation panels and integration panels in practice.

The second item which I wanted to share was that it may be possible, for example, to specify some circumstances, but I'm not sure if it is easy to predict all possible circumstances. So, I'm not sure whether it is a straightforward exercise to develop a complete and closed set of these specifications. So, just for information. And based on that, I guess I was going to ask whether the current status is of—that stability clause is useful, or would you still want to go back and update the LGR procedure? Thank you.

DONNA AUSTIN: Thanks, Sarmad. We weren't really sure, to be honest, whether the recommendations that we have from this working group actually have any jurisdiction, for want of a better word, over the LGR procedure. But we did think that if—and I take your point that

---

you can't always predict the circumstances, but we felt it was important to try to capture that it would be helpful to understand what the limited circumstances could be that could result in an update to the root zone LGR. So just to provide some predictability for registry operators.

So we thought this was important. We weren't sure whether this group actually could make this recommendation for the root zone LGR. And that ties back into one of the questions that we will get to later. So I think we kind of want this recommendation to stand. And then also understand whether the LGR procedure has an obligation to take that into account. So I guess that's a little bit of context for recommendation 1.6, if that makes any sense.

I see that Jeff had concerns. Or not concerns, but he's recommending that we don't need the words "extremely strong" in 1.5. If there's a presumption—is a presumption. And I'm not a lawyer, so I don't know about this question of rebuttal. But Jeff, go ahead.

JEFFREY NEUMAN:

Yeah, so, there's two kinds of presumptions, legally speaking. There's a rebuttable presumption, or there's an irrebuttable presumption. If something is rebuttable, obviously, that means that there can be exceptions to the rule. You would just need to indicate what those are.

So, what are we saying here? Are we trying to say that there can be reasons why it will not be grandfathered, or are we saying there shouldn't be any reason? I think we just need to pick one.

---

And then if there are circumstances where this presumption could be rebutted, then we need to kind of think about that as well.

DONNA AUSTIN:

Thanks, Jeff. And we went back and forth on this recommendation a little bit as well, and it kind of comes back to the question that we have below about what jurisdiction, for want of the better word, this group has over what the LGR procedure comes up with. Because we understand that if there isn't backward compatibility, then we're recommending that through a public comment process, this is called out and certain steps proceed.

But at the conclusion of that public comment process, it will be the generation panel that will decide—different question. At the end of that public comment period, will it be the generation panel that makes a decision about grandfathering, or is this recommendation 1.5 [inaudible] and the generation panel cannot [not] grandfather, regardless of what's in the public comment period?

So this is where it got a little bit tricky for us, because we weren't confident that we understood the process well enough. So I don't know if that helps or muddies the water a little bit. Sarmad.

SARMAD HUSSAIN:

I would just like to add that talking about a particular label and its status in the root zone would be out of scope for a generation panel's work, as it is at least currently described by the root zone LGR procedure. Thanks.

---

DONNA AUSTIN: Sorry, was that a question or a response to ...?

SARMAD HUSSAIN: Just a comment, I guess, that if generation panels—the way at least they're designed, or the responsibility which is given to them through the root zone LGR procedure, they normally work on the root zone LGR design. They don't work, specifically, for example, for certain labels as such. Thank you.

DONNA AUSTIN: Okay. Jeff.

JEFFREY NEUMAN: So I'm now a little bit more confused.

DONNA AUSTIN: This is why we struggled with this.

JEFFREY NEUMAN: Yeah. So we can still provide guidance to the label—or regardless of what the label generation rules come up with, we can still set a policy that if there's a change or—that anything that's already been implemented is grandfathered. We don't want to undo something that's already been done, right? So my question was, when you say something's a presumption, first of all, extremely strong doesn't really add anything. It's either presumption or it's not. But then when you say something's a presumption, it means

---

that it can be overridden by something. And that was the root of my question as to what could override that presumption.

DONNA AUSTIN:

Right, so I think the original wording we had was that existing gTLDs must be grandfathered. And then we stepped away from that a little bit because we weren't confident about making that policy recommendation because we thought the generation panel might have the last say. But I think maybe—because with the idea of presumption, I think you've put some doubt in my mind whether we want to leave that door open. So if it is the thinking of the working group that existing gTLDs must be grandfathered, then that should be the recommendation. There shouldn't be any—I don't know—wiggle room with that said. Jeff, and then Hadia.

JEFFREY NEUMAN:

I agree with that. And I don't think contractually a label generation panel can have any impact on the contract of an existing gTLD. And I don't think we're recommending—So while it's true that a label generation panel will have the ultimate say as to what's in the script, I think it's a very different question as to whether a decision by a label generation panel can be binding on a contract that it's not a party to, unless the contracts are changed for the future where this is put in. And I'm not sure why any contracted party would want to give a label generation panel—that it has no insight into or control over—any power to change something that a registry has already done. So I am much more in favor of the absolute rule that all existing gTLDs are grandfathered. And if not, if there's a reason why we can think of where there's an exception

---

to that rule, then we need to think of a process to deal with that exception. Otherwise, yeah. Thanks.

DONNA AUSTIN: Okay. Thanks, Jeff. Hadia.

HADIA ELMINIAWI: Thank you, Donna. So I do agree with Jeff. And I think what we are all saying is that the default is actually grandfathering those gTLDs, but we are also saying that the default is grandfathering but there could be reasons for not doing that. And again, as Jeff said, if we are saying that there could be reasons for not grandfathering those TLDs, then we should either give examples to those reasons, define a criteria, define a process through which this can happen, but we can't just like leave it open like this. Thank you.

DONNA AUSTIN: Thanks, Hadia. So kind of memories coming back a little bit here. So I think that was a reason for recommendation 1.6, with those, "specify the limited circumstances" because I think on the call that we had, Dennis raised the point that there could be limited circumstances associated with changes to the IDNA 2008 or the Unicode, if I have that correct. So I think that was what recommendation 1.6 was about.

And then maybe when we were doing that, we were thinking, well, maybe that affects 1.5. But I think what I'm hearing is support for 1.6, those limited circumstances should be specified. But for 1.5,

---

we should go with the absolute. Justine. And Edmon, did I see your hand up and you took it back down?

EDMON CHUNG: Yeah, Hadia basically said what I want to say.

DONNA AUSTIN: Okay, great, thanks, Edmon. Justine, go ahead.

JUSTINE CHEW: Thanks, Donna. I'd like to point out that we shouldn't be necessarily looking at each recommendation in isolation. I prefer to take the approach that all these are related. So if you look at recommendation 1.4, we are already asking the GP to call out reasons why they are making something not fully backward compatible.

So in my mind, from there, we would be able to tease out things that someone could evaluate to see whether it's big enough an issue for us to not grandfather something. And I suspect we might need to reword recommendation 1.5 to actually say what we want to say, which I think I'm hearing from the group is that the default is to grandfather subject to something dire enough to warrant not grandfathering it. Thanks.

DONNA AUSTIN: Thanks, Justine. So I don't know if you've seen what Jeff has put in chat, which is—well, you can read what's in chat, and I think Jeff's going to speak to it. Anil.

ANIL JAIN: Thank you, Donna. Yes, I agree that most of the members agree for grandfathering. But my suggestion—I don't know whether it is implementable or not—there may be a possibility that a particular variant may not be suggested to be grandfathered. So what I suggest, in case it is possible, can we take a public opinion in case a situation comes when an existing variant is blocked or it is not supported by root zone LGR? Can we take a public opinion in that particular region before granting a grandfathered status? Thank you.

DONNA AUSTIN: Thanks, Anil. So I think your suggestion might be covered within recommendation 1.4, which talks about public comment process. But if you could just review that and see if that's the case. Jeff, go ahead.

JEFFREY NEUMAN: Thanks. So I agree with Justine that we should read all these together. And there should be a default. But when this is implemented, this is going to be implemented in contracts. And I don't think we should make a label generation panel have any kind of power over contracts to which they're not a party to.

So I have no problem with the label generation panel making recommendations and not being subject to a bilateral discussion between registry and ICANN to figure out if it can't be grandfathered, what can be done. But again, we have to think about this in terms of implementation and contracts, and the label

---

generation panel is not a party to the contract. So I still maintain that we can say except as set forth in whatever exceptions, all gTLDs and their delegated and allocatable variant labels affected by aforementioned exceptions will be grandfathered. But then we just need to deal with a process that involves ICANN and the registry, because it will affect their contract, not to mention the effect it could have on end users. Thanks.

DONNA AUSTIN: Thanks, Jeff. Justine.

JUSTINE CHEW: Thanks, Donna. Just very quickly. I think what Anil meant is what Jeff said. I don't think he was talking about the public comment where the generation panel proposes something that causes the problem. I suspect he's talking about after that public comment closes, and we see what the generation panel has to say, then somebody else—or there should be another process, which Jeff is suggesting that too, take place in order to see the ramifications and whether we want to then have something grandfathered or not grandfathered. I'm not saying we as in the EPDP team, but there should be a process. And it makes sense that it should be between ICANN and the contracted party. And ultimately, the decision of the Board, I guess. So maybe we might want to think about working something along those lines in if the team is agreeable to that. Thanks.

---

DONNA AUSTIN:

Thanks, Justine. Unfortunately, we're at time. We're actually one minute over. But I'm pleased we had the opportunity to have this discussion. I think what I'd like to do, based on the conversation we've had here, we'll try to do an update to the recommendations. I think certainly, we can change 1.5 to a "must be grandfathered except as set forth in 1.6." And I think we need to maybe tidy up 1.6 a little bit.

I think from the leadership team perspective, and when we had this discussion earlier within the group, there was some concern about what the GP could do and what the working group could do. But I think I'm coming to the conclusion that we just need to make the recommendations that we can do within our bounds, and we will try to sort out the implications of that on other processes. And Sarmad's helpful, because he can point us in the right direction. So, let's not worry too much about that. Let's just get our thinking right on what we want to see in these recommendations.

Okay. Yeah, I agree. Justine's saying that the discussion we've had today may have already answered some of the questions that we had at the end of this document. So that's helpful as well.

So we're going to call it at time. We'll try to do a quick turnaround on this document and get a revised version out to the list, and then hopefully, we can make some progress on drawing a line under A6 as well.

So thanks again, everybody. I think we've had some good conversation and we're making progress. All right. Enjoy the rest of your day, evening, wherever you may be. Bye all.

DEVAN REED: Thank you all for joining. Once again, the meeting is adjourned. I will end the recording and disconnect all remaining lines. Have a wonderful rest of your day.

**[END OF TRANSCRIPTION]**