ICANN85 Mumbai | CF – GNSO: CPH TechOps Work Session (2 of 2)
Sunday, March 8, 2026 – 10:30 to 12:00 IST

DEVAN REED

Hello and welcome to the CPH TechOps session 2 of 2. Please note this session is being recorded and is governed by the ICANN Expected Standards of Behavior, the ICANN Community Anti-Harassment Policy, and the ICANN Community Participant Code of Conduct concerning Statements of Interest.

Please observe the following guidelines to participate in this session. I will also post them in chat for your reference. Only questions posted in Zoom chat are identified as a question, will be read aloud during the session as time permits when directed by chair of the session. If you wish to speak, please raise your hand in Zoom or otherwise as directed. When speaking, please state your name for the record and speak clearly at a moderate pace. Thank you and I will now hand the floor over to Roger and Mark.

ROGER CARNEY

Thanks, Devan. All right. I think we're going to go ahead and jump in and let Michael get started on his IDN and update. So, go ahead, Michael.

MICHAEL BAULAND

Thanks. Michael here. I would follow Peter's example and take a step back and provide some more basic fundamental information

about all these IDN-Related Groups, what's this about, why is it important for us as registries, registrars, and not go so much into detail about the upcoming EPP extension we are developing, but more provide some background information.

Next slide, please. So, what are variants? These are actually defined via LGRs, these label generation rules. This is a standard RFC 7940. It's a way to define what code points are allowed in a label and also what are the variant relationships between code points. I have here an example from the Latin LGR, and you can see here that the first line just shows a single code point, O12F, which just means that code point is a valid code point in any label if the Latin LGR is used, but it has no variant relationships.

The second code point, O131, which is a Turkish dotless "I", as you can see, there are several variant entries for that. It even goes longer than this, but I cut it off. You also see that most of the variant entries are of the type blocked. So, these are variants, but nobody will be able to have a label with these code points replacing the O131, but still a registry running this need to make sure that labels having these replacements are not available for anybody else if one of the, well, not just anybody else, for anybody actually, if that label exists.

There's also in the first line you see O069, which is a type dotted. This is later in the LGR, there are those rules, actions that define how different variant types are used, and the dotted one is actually an allocated variant or allocatable variant. So, in this case, if you

start with a dotless “I”, the standard I, O069, is an allocatable variant. As an example, I put in here Turkish city town, Aydin, I'm not sure if the pronunciation is correct, but that has the allocatable variant of Aydin with a standard I. And the other allocatable variant relationship in the Latin table is the sharp S, SZ, like in Straße, that has a variant with a 2S.

Yeah, next slide, please. So, as I said, there are lots of block variants, and there's this ICANN LGR tool, it's publicly available for anybody, just log into your ICANN account, and you can validate any label using any of the reference LGRs.

You can also upload your own LGR, which you have created to check that, and you can type in any label, like here, Straße in Aydin, which means like a street in the city of Aydin. And if you validate that label, the tool will say, oh, too many variants, which doesn't mean it cannot calculate it, it just means it cannot or will not display those labels right away, but it will generate them in a background process, and you can download the file containing the labels.

In this example, this domain has 9,216 variant labels, including the label itself, but only three of those other labels are allocatable labels, all the others are blocked. Most of them got introduced via transitivity from other scripts.

Yeah, next slide, please. So, if we take a look at the second level in just an example TLD, which is using the reference Latin LGR table,

the set of allocatable labels, actually depends on with which label you start.

So, if you start with Straße in Aydin with a sharp S and the Turkish dotless I, then you have kind of all three combinations as allocatable variants, exchange the sharp S with double S, and or exchange the dotless I by a standard I. If, however, you start with a label where you already have the standard I, then there is only one allocatable label, because the variant relation of both of these is just allocatable in one direction and it's blocked in the other. So, we can't come back from a standard I to a dotless I as allocatable variant. So, the primary domain actually matters.

Next slide, please. And in the IDN PDP, the policy was created to make sure that for variants, there's no user confusion, and this requires the same entity set. So, all allocatable variants, because block variants are anyway not available to anybody, have to be registered via the same registrar. This needs to be checked by the registry, because of course the registrar cannot check it. The registry has to ensure that all allocatable variants, if created, this is done via the same registrar, and they also need to belong to the same registrant, and that needs to be checked by the registrar.

Next slide, please. And what about Latin diacritics? This is a current ongoing PDP, and this basically deals with issues like Cafe and Café, one with a standard E and one with an accent, which are not variants of each other according to the reference LGR, but still are likely confusingly similar.

So, if someone wants to have both of these TLDs, under the current policies, it's highly unlikely that they will get both, because they would be very likely considered to be confusingly similar, and therefore only one could exist. And the LDP DP tries to create or is creating policies that make it possible for both TLDs to coexist. And we create the rules similar to variants, so for all technical purposes, for example for EPP, it really makes no difference whether you're dealing with variant TLDs or whether you're dealing with TLDs which are combined via Latin Diacritics set.

So, once those TLDs are in the root zone and available for all practical means, it's no difference between variants and the LD set. That is also one of our main goals in the LDP DP, to not make another special case where registries, registrars have to do special things, because that would really make adoption difficult, but to do it exactly the same way, which makes it easier for all parties.

Next slide, please. So, most of you know about the RSP, Registry Service Provider Evaluation Program, and there was one question about what IDN level you want to support. And there are basically three IDN levels that registries or registry service providers were able to choose from.

I'll quickly go through those and say what the difference is. So, IDN level one, you need to do this if you want to offer IDN domain names. But this is the most-simple case. You do not need to have support for variants. There's no special EPP extension required for this. The only thing that you need to do is you need to block

variants. Independently of whether these are allocatable variants or block variants, you need to make sure that if one domain is registered, all other existing variants of the set are blocked.

Next slide, please. So, in the example, we registered this first domain name. Then all of the three allocatable variants also need to be blocked. Then, of course, the other thousands of blocked variants, which we talked about before, always need to be blocked.

Coming to level two, this means that you as an RSP do want to support variants at the second level. But any registry that want to have variants available for the registrars need to implement level two. And for this, some EPP extensions are necessary and registries may allocate or block variants. Next slide, please. So, if you start with the first one here, register it, then this becomes a primary. As we saw before, it matters what the primary is.

Next slide, please. And the other three allocatable variants are available, but only to the same entity. Everybody else trying to get those will see that they are not available. Next slide, please. If you, however, start with this one, if a registrant via registrar registers this domain name.

Next slide, please. Then the consequence will be that the first two need to be blocked, unavailable for anybody. And the fourth domain name with only ASCII characters that is still allocatable to the same entity.

Next slide, please. Now we come to level three, IDN support. This is registries that want to support variants, not just at the second level, but also at the top level. So, it means they will be able to offer variant TLDs or LD set TLDs. Again, a special EPP extension is necessary to implement this, and registries may allocate or block variants. Next slide, please. So, here we are seeing the LD set of .cafe and .café. And if someone registers this first domain name, Straße in IDN, with a sharp S and a Turkish dotless I under the ASCII TLD .cafe, then this became the primary domain name.

Next slide, please. And all other seven domains become allocatable again for the same entity. I didn't put it here because the slide was too narrow for this, but it's always the same entity here.

And next slide, please. If you register this domain name, question to the group, how many domains do you think will be allocatable for the same entity? Any guesses? The thing is at the ASCII TLD .cafe, we have the same situation which we saw before on level two. The first two ones are blocked, and the ASCII only second level domain is allocatable. But on the other TLD .café and the IDN, we do not have a primary label yet. And at this point in time, all four domains are still allocatable.

Next slide, please. It really depends on the first domain in that TLD that gets registered, which then defines the disposition value of the domain names in that TLD. So, if we then register this IDN with .lsi, then all the other three of that still are allocatable. And at that

point in time, we have two primary domain names in this set. Each TLD will have their own primary domain name.

This is the situation as before, but if we then the first domain registered under the IDN TLD is the ASCII only second level domain name, then all three other previously allocatable domains become blocked in that domain name. And the primary domain in the IDN TLD is the ASCII only second level domain. Questions so far? Yeah, Jan? Rick. I mean.

RICK WILHELM: Some of that's a little counterintuitive. I mean, I believe you because you're you, but, more dangerous is believing me, right? But some of that is a little counterintuitive because it's sort of like it's almost quantum in that like the TLDs are linked, but they're not. And depending on how they're the order of events, this is capital H hard, right? It's just an observation.

And I don't doubt a word of what you're saying. But and I'm sure that you have this implemented in code. But I hope that the top of that file is got one of those headers on it that says, if you think you understand this, you don't, right? You know, one of those. I mean, we've all seen that. Oh, those of us that write or have written code have seen it. I hope that that file has that header on it. Thank you.

MICHAEL BAULAND: Thanks, Rick, for the kind words. And yeah, this might be counterintuitive, but it's actually necessary because those two

TLDs could actually have different LGRs. They need to be harmonized. They need to be consistent, but they may be different.

One example would be that under the IDN TLD, you allow all Latin script domains at the second level, whereas at the ASCII TLD, you may possibly only allow ASCII domains at the second level. That might be a sensible use case. And so, for that reason, you need to have, take a look at all registrations in each domain kind of separately. You just need to make sure that its same entity set that spans all TLDs. So the same entity always needs to be the same across all TLDs, but the variants are, whether they are allocatable or blocked, that question is only within a TLD. Yeah, question?

MACIEK PIASECKI Maciek Piasecki here with a little bit of follow-up to that. I was actually the reviewer and pen holder for the public comment by ALAC to the Diacritic PDP. And I'm wondering what are the cutoff points here where you consider a pair of characters confusing or a variant of each other? Because I imagine this is not done on a case-by-case basis, since there are apparently a thousand or so Diacritic characters. So, is there some code for that? And is it constantly being debugged? And is this code open source? Can we probe it a bit?

MICHAEL BAULAND Okay, thanks. Good question. So, for variants, this is a clear cut, clearly defined case. The LGR clearly says what is a variant and

what is not a variant. Of course, registries at the second level can create their own LGRs. They can have different definitions of what a variant is. They can even make this E and ‘E with an accent as variant if they want to do that at the second level. But this is clearly defined.

What is confusingly, that's not really defined. This is subjective. There's no possibility to objectively say something is confusingly. Some people might find it confusing. Others might not find it confusing. So, it's always subjective to some extent.

Regarding the LD, Latin Diacritic set, this is also clearly defined. We look at the Unicode definition. Unicode has a clear definition of what modifying characters are called Diacritics. And the characters which are in Unicode defined as a composition of an ASCII character plus one or more Diacritic characters, those are diacritic characters and are in scope of the LD set.

Okay. Next slide, please. So, with this, I only quickly want to talk about the new EPP extension which we are currently working on that covers all the technical requirements, implications that come with these new IDN level one, two, three things. I won't go into much detail of how we do this, just a brief overview of what this extension is intended to do.

So, the same entity set is created when the first domain is created. This is like we saw in those examples, as soon as one of the domains is actually created, allocated, this creates the same entity set because from that point on, the registry needs to make sure that all

elements of that set are either blocked or reserved, withheld for this same entity.

And then we have that those domains which are allocated can be updated to the status of being allocated, so that they are active or they are within the registry. And they can also be updated to switch back from being allocated to be allocatable again. So, this is basically creation and deletion of those same entity domains.

Next slide, please. And then for check info transfer, we have the basic rules that check will tell you whether a domain is part of an existing same entity set. So, if a same entity set has already been created and you do a check command of one of the domains within that set, the registry will have to behave depending on firstly, does the registrar understand same entity sets? If not, they will just say, no, it's unavailable. They could put in the reason of the available. It's part of the same entity set, but that's up to them.

If the registrar supports the same entity set, it needs to check whether the registrar is the sponsoring registrar of the existing domain. If yes, it will tell the registrar, yeah, it's available, but with an extension telling him this is part of an existing same entity set and this is a primary domain name. If the registrar is not the sponsoring registrar of the existing one, it will also say it's not available, possibly with the information what the primary domain name is via an extension.

In the info request, again, assuming the registrar understands same entity sets, if not, the info will just return the standard info response

as if no same entity set exists. If the domain exists, you get the information of the domain. If the domain is part of a set, but it doesn't exist, you don't get any information because it's not existing.

If the registrar supports the same entity set, the info will return all active domains in the same entity set and also all primary domains of the same entity set. This is important. It's not and never will be, registry will never return all potential variants of the same entity or all potential elements of the same entity set because as we've seen in the example Straße in Aydin, this may be thousands of variants even in this rather simple case and it can go up to millions or billions of variants.

It will only ever return the really active domains and that should be a rather small set of domains because registries are also required by the Aydin PDP to have the maximum number of allocated variants slow. They should not allow allocation of thousands of variants.

And then finally, there's a change to the transfer command because that needs to make sure that all variants, I shouldn't say variants because we extended this to the same entity set to include non-variants like LD set or other use cases, but most common use case in our case will be variants. So, a transfer will need to transfer all of the existing active domains in the same entity together with the main domain name.

This also means that a registrar sending a transfer command but does not support same entity set, they should not be able to do the transfer at all because otherwise they expect to transfer a single domain and might suddenly have like 10 domains in their portfolio which they have not been aware of. So, the transfer of more than one domain should only be possible if the registrar understands what they are doing and they can prove this via the extension. And I think that's it. Any questions, comments? Yeah.

ROGER CARNEY: This is Roger. What happens on renewal if you are not renewing the primary but you wanted the allocated ones?

MICHAEL BAULAND: Good question. So, this is a bit up to the registry policy. They can have different life cycles for all domains of the same entity set but they could also have a combined life cycle. Say only the primary domain name has their own expiration date and all activated variants just follow in suit or they can have their own expiration date and once they get activated, they get their own date. But expiry in itself doesn't really exist in gTLDs. It's always auto-renewed.

So, you would have to actively delete domain and there it's also up to the policy of the registry. Either they forbid the deletion of the primary domain if there are still variants existing or they can have a policy that if you delete the primary all others will be deleted as

well. But it must not happen that the primary is deleted and the others in the same TLD remain existing because then you have problems with the disposition value as possible.

You have no primary anymore. So, primary may only be deleted if none of the others exist and whether the registry wants to delete the others or forbid the deletion of primary is up to the policy. We won't prescribe that.

ROGER CARNEY

So, maybe a question for the registries. I mean, is there thought that a primary could be reassigned and then obviously that also changes the allocatable variance under that. So, if you have a primary that someone no longer wants but they still want two of the allocated ones and they want to switch the primary to one of them, is that something? I mean, Rick said it doesn't seem really logical necessarily sometimes here and it's like that just adds another exponential problem on that that seems like registries would probably do a package instead. But just thinking out loud.

MICHAEL BAULAND

I'm a registrar. So, he asked the registries. But yeah, we are also an RSP and actually we hadn't thought about whether we want to make that possible or not. But the case is a sensible one. So, it could switch to slide 22 quickly, if that's possible.

RICK WILHELM

I mean, I'll offer. I think that a lot of these corner cases are super esoteric and theoretical, right? And a registry implementing this stuff is probably going to do things that are simple because a registrar selling this stuff won't want to have all sorts of fiddly bits, technical term, on the other end in their shopping cart and shopping experience.

And they're not going to let registrants do sort of weird fiddly unbundling with these variants because it's going to be hard enough. And so, they're going to do something where if you dip in and you grab one of these and the other variants sort of come along, then you're going to get that package and you're going to have that bundle and you're just going to have that bundle.

And if you want to do something on your own to switch between what you want to be the primary one in your mind is where you direct your traffic, go right ahead. But you're not going to do some, they're not going to the registry and the registrar aren't going to price it such that, oh, you might save a euro by, oh, only doing two out of the three variants. You're going to just pay because these are going to be expensive for the registry and the registrar to sell and service.

This is not going to be a five-euro domain. These are going to be, there you go. Yeah, they're going to be expensive. So, all that sort of fiddling about is just not, that's my prediction. If I'm the product manager, I'm not going to put that many options because the

customer is not going to be sophisticated enough to deal with them.

ROGER CARNEY

Yeah, and I think that last part is the interesting part because I think a registrar will have a hard time explaining this to an end customer to begin with. But I would be careful about some registrars not wanting to do it because if it's a dollar, they're going to keep the dollar. And if they can break a bundle and still make money instead of making no money, some registrars may try to do that. So, I would encourage our registry people to make some good policies around this so that they can't do certain things. So, yeah, thank you.

MICHAEL BAULAND

Coming back to your question, this would be an example where you might want to change the primary because if you started with the wrong one, so to say, there's no way you can ever activate the top IDN domain name. You would need to change the primary. But I also think, as Rick said, this is quite a corner case and this is only necessary for asymmetric relationships like we have in the Latin case.

Many LGRs don't have this. They are symmetric, then you don't have that problem. So, my guess is that there won't be any technical features to make that possible because it will make the whole process quite complex. But there might be a support call

asking them, could you maybe switch that for us? And then some geek with an admin interface somehow changes and switches the primary and makes this available for the registrar if they ask nicely.

MACIEK PIASECKI: So, did you do any test runs with the registrars if it's feasible for them or whether they understand this system and they are able to provide relevant customer feedback to the customers rather than just saying computer says no?

MICHAEL BAULAND: Yeah, this is all still a work in progress. This whole EPP extension is currently being drafted in the IETF. If you follow the main mailing list, you can see it and you can join in to the discussions. And if you think it's necessary to have something included in the EPP extension to change the primary domain name, voice your opinion on the list. Yeah, there's a public discussion. It's not final yet. And therefore, there also haven't been any test runs also yet. It's still a work in progress.

ROGER CARNEY: Great. Thanks, Michael. And just my last comment, you make the example and I'm going to maybe make Rick cringe here, maybe simple at the TLD level because the Latin Diacritic, I assume, can be more than two TLDs?

MICHAEL BAULAND Yeah, both variants and Latin diacritics could contain more than one variant TLD.

ROGER CARNEY So, that same entity could have six of these and then the blocking and allocating across them have to be managed, right? Yeah.

MICHAEL BAULAND But again, it's a highly unlikely case. I think even having variant TLDs in the next round, I don't think we have like hundreds or thousands of variant TLDs. My personal opinion is we have maybe 10 or 20 variant TLDs. I don't know, maybe 50. And I would say 99.9% of those variant TLDs have just two labels. But that's just my gut feeling. I have no information.

ROGER CARNEY And speaking of the Next Round, where do these things stand? As you said, this is still work in progress for Latin diacritics. Is the next round going to implement the first IDN PDP?

MICHAEL BAULAND Yeah, the IDN PDP is definitely part of the Next Round. That was a requirement. That's also why you had to define your IDN levels one, two, three in the RSP evaluation, whether the LD PDP is in there. It's highly unlikely. It's not totally impossible. But technically, it doesn't make a difference because whether it's there

or not, it will be the exact same PDP extension and the behavior will be exactly the same.

ROGER CARNEY

Great. Thanks, Michael. Any other questions for Michael? Again, a great, I guess, background and step back to look at how these things are potentially going to work. So, if there's no other questions, okay. I think then we can jump on to an update from Rick on the RDAP profile.

RICK WILHELM

Very good. Thanks. So, we're going to talk a little bit about the RDAP profile. You can flip to the next slide.

RDAP, of course, has been in production for a while. We've got the dates here on the slide. We've been working with the version of the profile that was done in 2024. It's hard to believe that it's been over two years that the profile was finished. Good grief.

And of course, we were working on it for a while before that. This thing started before the pandemic, if I'm pretty sure, right? Like we started working on this thing because this is a second version of it. So, anyway, and of course, this is all, yeah, the 2019 version that we had is obsolete. And so, we did an update. We've got a February 2024 version.

So this is, there are a lot of people that have been active in the community for a while that probably don't have memories from when there wasn't RDAP, which is making me shudder. So, enough

of that. Okay. So, and there's links there and such. So, we can go ahead and flip on. But the profile itself, of course, has got two documents. And this was on the prior slide, which I didn't read. But it's the, actually, can you flip back one slide?

Sorry. It's got the Technical Implementation Guide, aka the TIG, for those that traffic in acronyms, and the response profile. And so, the way, and when this was built, originally conceived, probably back in like 2017, when we had nothing but unplowed ground in front of us, we thought that the idea would be that we would separate the technical bits away from the policy bits, such that the policy bits could change. Because at that point, we didn't really know how much of the policy bits were going to need to change over time, right?

Things were still, things were still very much up in the air around some policy stuff. The registration data policy, the PDP was not yet starting. And so, at that point, most registries were thick, and the, but not all registries were thick. GDPR was still a controversy. And so, if you put yourself back in that, but CommonNet were still thin, and was actually under an exception or an extension about going thick.

And it's hard for a lot of people to remember that that's sort of the way that the world was. And there was CommonNet were operating under an exception for the policy for various reasons, set that aside. But GDPR was coming into being. And so, we knew that

the technical people knew that some of this policy stuff was in flux, but it was not at all clear how things are going to come down.

And so, what we did is we separated the documents into two parts, the technical implementation guide, which was supposed to be mechanism, and the response profile, which is supposed to be policy. And so, this hasn't been said on the record in a while, but that's for those of you that are looking and doing RDAP stuff and see these two documents, and you're cursing or muttering or otherwise frowning at why you're dealing with two documents.

This is why? Because back almost 10 years ago, the future was not yet written, perhaps obviously. And it was not at all clear how this is going to turn out. Okay. So, now, Zoe, please flip to the next document. So, one of the things that's happened relatively recently is the ICANN team built the ICANN RDAP conformance tool, which is out there, and the URL is at the bottom.

And what it does is it tests whether a particular RDAP implementation of a server service is conformant to the RDAP profile documents. And so, it's a test suite, essentially. And so, the, and ICANN built this in order to help it determine whether or not people are fully complying with the contractual requirements.

The contractual requirements being, one, you have to provide an, if you're a registrar or registrar, you have to provide RDAP, an RDAP service. It has to be performant. It has to meet certain SLAs about how fast it works. But then it also has to do the right thing.

And so, it's set to give back the correct answers, but then those answers have to be correct in both content so that when you query foo.tld, it has to give back the answer for foo.tld and not bar.tld because that would be incorrect. But then that, the specification of that has to come back in the way that the RDAP profile says it does. It can't be gobbledygook because otherwise a client that's querying that would be utterly confused about what it says.

And that seems simple enough because, but those of you that have done software testing know making sure that software fits to the spec, you sometimes find bugs in the code and then you sometimes find bugs in the spec. And sometimes it's unclear whether you're finding bugs in the code and bugs in the spec and you have a meeting and you hash these things out.

Sometimes they end in laughter, sometimes they end in tears and occasionally they end in departures from the organization. Hopefully they don't end in arrest or charges, right? Or lawsuits. So, the ICANN team built this and it was not an overnight exercise. So, it was a lot of capital R, capital W real work. And so, in so doing, they built this thing and then they were testing the suite against everybody's implementation and sort of saying, you know, hey, Rick, your RDAP thing, we, we put in a query for foo.org and it gave back an answer for bar.org. And I would say, oops, that's a bug on my part. Let me fix that.

And then they might say something else. I'd be like, no, I think that's a bug in your tool. And Andy would say, oh yeah, that's a bug

in our tool. And we'd kind of go back and forth. And then sometimes we'd get to an impasse over this, and this is all hypothetical, but just an example. Don't let the truth get in the way of a good story, as it's often said by our uncles.

Sometimes we'd be like, well, there's kind of an ambiguity in the spec here, right? And so, what happens is that there's a couple of times where folks have come up with things. The folks at Tucows have sent a couple of items to the list.

Sarah, we talked about these at ICANN84. This is in the zoom recording where there are two issues that they ran into, unsurprisingly, because ICANN was doing this full suite of running the full spec and going over everybody's stuff. They ran into a bunch more stuff, right? But I'm guessing that some of what the Tucows team came about with dialogue, either they were testing their own stuff or they were talking to ICANN.

So, anyway, let's go to the next slide. So, one of the things that when ICANN team was creating this thing, the stuff that they kind of came up with, you can see what here is, had sits on the screen. They kind of said, you know, we might need to fix some things that were like some changes that -- where we've always said that when we're doing the RDAP profile, we don't want to establish policy. And if you'd go back through all the -- we want to memorialize the implementation of policy, but we don't want to establish policy.

And if you go back through the transcripts and run the AI across all that stuff, you probably would find instances of me saying that

probably a hundred times. If like to the point of where I would just say that constantly, like we are implementing policy. We're not setting policy. We're taking direction from the policy team.

And so, one of the things that they identified was places where our implementation does not precisely match with the policy. There's just a little bit of a difference there. And so, like nudging those things exactly into places. There's a couple of places where the standards have advanced since our original implementation. And we've just got to do this. There'll be some suggested updates to these standards.

Maybe there were RFCs that have been updated since then, or places where we cited internet drafts that have turned into RFCs or some other things. And then there's some spots where the language is a bit ambiguous and it either allows for multiple interpretations or the only way to get it to fit is to interpret it in a certain obtuse way. Like, as it's been sometimes said, it depends on the definition of infamously. So, that would be following the category of clarifications.

So, let's flip to the next slide. So, then in addition to some of those things, some of the things that the ICANN team are suggesting is, when we revise this thing, there's some things that we might possibly consider. One is suggesting now that we have the benefit of pushing almost 10 years of experience from when this work originally started, we can probably consolidate this into a single document because it's clear where the policy is coming down on.

The policy is no longer wiggling. And so, this notion of the technical versus the technical implementation versus the policy guide, we can put it into one document, is probably okay because the policy is not going to be wiggling about and such.

Secondly, there's probably some places where once it's put together under a single spot, being able to restructure it to put all of the discussion about a certain topic in centralized spots. Third is create a mechanism right now for doing minor changes for errata or simple updates to make it be easier to patch the document without doing a full revision. Right now, our process for doing a full revision involves the CPH and ICANN getting together to discuss it and request a new version. And it's a pretty heavyweight process.

They're suggesting it would be a good to include a mechanism to allow us to do small updates that are non-controversial. But that mechanism would have to be invented. But basically, to do an errata mechanism that if we come across minor, what my words, syntactical changes that don't have semantic impact, right, that we could do those kinds of things. Maybe, and we would have to define what those things are.

Some other things, add more examples or informative text and further have a way to update those without it changing the semantic text, right? So, separate the informative text from the normative text. So, the informative text could be updated easily so that we could add examples in and help clarify the document.

The fifth one there, reconsider the source form as HTML versus PDF. Right now, the document is published in PDF form. Why HTML? Well, if as everybody here knows, and I would assume I'm safe in saying everybody is deep linking into a PDF document is extraordinarily difficult.

And so, whereas if you publish something as a PDF or as an HTML document, you can deep link into it and reference into it. And the IETF shows us that is clearly possible.

And then the last time is maybe a little bit harder to do, but a possibility of considering a public issue tracker or some sort of issue tracker so that if people do find things right now, our mailing list is the, which is not very heavily trafficked and certainly not public is the only mechanism we have for people to report issues as the Tucows folks did months ago. But if we had a public mechanism, whether that be a GitHub or something like that, where we could post issues, that that would be something to consider.

I think that's the last slide that I have Zoe. Yeah, okay. So, Andy was not abled. Andy Newton from ICANN, who's the principal author of the test suite was not able to be here tonight. Sorry, Eastern time because it's nighttime Eastern time.

And the ICANN team has a big long list TM of detailed items about things that they have suggested that deserve further discovery and research in addition to the two items that the Tucows team has as far as detail on. So, we have a good solid punch list that is meetings

worth of work to attack. So, I don't want to suggest that this is a, oh, you know, like, hypothetically, Roger and I would get together at a week for two days in Chicago and bang this thing out and propose something to people.

This is going to be a fair bit of work. Right. But I think it's worth discussion and such. So, let me stop there. I think Roger was first, and then Sarah.

ROGER CARNEY

Thanks, Rick. Yeah. And maybe we just solve it over beers tonight. So, I don't know. No, I agree. I think there's a lot of work here that the group should get back to working on. Just two comments on the consolidate to a single document because policy set, I would be very cautious of that. Mostly because there's current policy work going on the privacy proxy stuff that may impact this. The discussions coming up on accuracy may impact this as well. So, I wouldn't say policy is set on this.

Now, do we know more than we did 10 years ago? For sure. I still think the two documents seem to make sense, but again, obviously up for discussion, just an opinion there. And on the public issue tracker, I liked the idea of viewable here, but I don't want people to be able to post issues to a tracker that an RDAP Working Group would have to look at.

I think that that has to be resolved at the ICANN level. And then ICANN could create something that says, okay. And then Tucows,

either one of them that wants to look at it can look at it and see how it's going. But I would say not a public posting, but a public watching maybe something is a good idea, but just my comments and I will turn this to Sarah.

RICK WILHELM Yeah, I think both of those have points of merit. Yeah, thanks.

SARAH WYLD Thank you. Hi, this is Sarah Wyld from Tucows. I want to talk about my issues for a moment. So, as Rick referred to the profile has a really important piece to it, which is that it says that if there's a conflict with the policy, then the policy is what overrides. And we've run into that twice, Tucows. But the one that's specifically in my mind right now is about policy number two point, sorry, profile item 2.2, the registry domain ID.

So, ICANN compliance has a concern with my RDAP implementation that has been open since I think November, long time. And they're really insistent that we, the registrar must publish or publish and redact the registry domain ID, but we don't collect that data. So, I couldn't publish it even if I wanted to. And we've had three and four back and forth with compliance where I'm saying the policy says, if you collect it, then you publish it.

The profile says you must publish, but the policy is an if, so there's a mismatch and ICANN compliance will not accept that there is a mismatch. So, that's where I'm stuck is, I love the idea of a public

issues tracker. I love having a mechanism to fix errors, but if compliance can't understand that there's a mismatch, I'm not sure how to say it any more clearly. And so, maybe somebody has ideas or maybe we can just get it updated. Thank you.

RICK WILHELM Sure. So, compliance isn't accepting of the line out of the profile that says, is that there's a conflict?

SARAH WYLD No, they're not accepting that. I've said to them, like I've quoted, here's what the policy says. Here's what the profile says. Look, it's the same words. It's not matching. And they've said, like, we do not see a mismatch.

RICK WILHELM But the response profile says, in cases of conflict, the policy wins. And they're just saying,

ROGER CARNEY They're saying there's no conflict.

SARAH WYLD I just pasted it in the chat. Based on the requirements, ICANN has not identified a conflict between the policy and the profile as it pertains to the registry domain ID. That's what they're saying.

ROGER CARNEY And I think, Sarah, this is maybe something we can bring up in our compliance.

SARAH WYLD It's on the list.

ROGER CARNEY Yeah. Because I think that this is bigger than Tucows. Yes. And maybe other people are doing it like ICANN wants them to, whatever. But the fact is, if they can't see the conflict, we need to raise that in our compliance and say, well, just because you don't think you see it, it doesn't mean it doesn't exist. So yeah, thank you.

SARAH WYLD Thank you. And I would be interested if anybody else is having a similar concern or back and forth issue, or if anybody else is publishing this and they don't have to, like, I'd like to hear what you're doing.

RICK WILHELM Maybe that's something that gets brought up in the joint CPH meeting. Don't we have a joint CPH meeting at some point? Or did we already have that? Gordon?

GORDON DICK I was wondering if you'd had any engagement with the technical services team to understand their technical understanding of it, because obviously they feed into the compliance team and might help educate there.

SARAH WYLD Thank you. This is Sarah. We started, like, before we heard from compliance, I emailed technical services and said, hey, I think we have found a problem. And so, we're going to follow the policy. And they wrote back with a response that was like, okay, you do what you got to do.

ROGER CARNEY Please.

MACIEK PIASECKI Maciej Piasecki here. So, given that only after a few years, there's so much rework needing to be done. Do you think that the original PDP work did not address future proofing enough? Or maybe on the contrary, it means that the system is so robust that it's open to being reworked as the landscape changes dynamically?

RICK WILHELM Sure. Good question. So when I went through, and I purposely didn't put -- so, short answer, I'm still very comfortable with the work that the group did when the profile was created and then updated. So, I'm still very comfortable with it.

When Andy and his team have a chance to present their work to the group, I think that the group will find that the issues that they uncovered are interesting, right? Like, mmh, okay, yeah. There's a category of them that are like, oh, that's a bug. And then there's a category of them that, oh, that could have been better worded, right? But I found all of them, all is probably too strong, but the vast majority of them to be within bounds of reasonable expectation for what you would expect.

So, it's aged reasonably well. You know what I mean? You know, even food that you put in the freezer is going to suffer some freezer burn, right? But there's a shelf life that it has and it's hung in there. It's aged reasonably well. So, I'm okay with it, yeah. And if it had survived, if it had gone through and they'd written the conformance suite and it hadn't found anything, I probably would have been a little bit skeptical, right?

It's a little bit like when I test my own code and it works the first time, I almost don't like that either, right? I kind of want it to break because then it's like, okay, it got a good shaking, right? So yeah, it's okay, yeah. Is that fair enough? Yes.

ROGER CARNEY Thanks, Rick. Yeah. And I was going to say, maybe Rick probably paid you to say that. I'm not sure because that's a good advertisement. Because I think the group did do the job and actually made it at the time, again, 10 years ago, whatever it was

when it started, they knew changes were going to happen and it was purposely built for that.

And when you look at the RDAP Working Group, it's not a heavy-handed group. It's a pretty efficient group that gets through and updates it from the 2019 to the 2024 spec. There were some changes there and it was actually a very efficient change through the RDAP Working Group. And that Working Group was put together specifically because we know it's changing. So, I definitely think it's doing its job. So, thanks.

RICK WILHELM

All right. We'll flip to the next topic. This is the IETF update. So, we're going to go ahead and I'm going to present this stuff for Jim. Jim is otherwise engaged. So, he presented these slides and we've got until the top of the hour, right? Okay. So, I'm going to try and bang through this.

So, this is really a preview of the IETF meeting. The IETF meeting is coming up next week in Shenzhen. And so, Regex, I can't remember exactly what day it meets, but that's published in the IETF schedule. I should know by now, but I didn't check it because I'm working on ICANN stuff this week. So, here's a quick run-through of kind of what's on the agenda.

So, first thing recently, this is recently published RFC 9910. This is an update for RDAP to let you search for RIR things, right? So, to let you do “searching” across RDAP for IP blocks and AS numbers and

things like that. I'm pretty sure that this is, I don't have it, I didn't write it. I think this is Yazdeep [ph], one of Yazdeep's things that he's been doing RIR work there. So, that just got published.

The next one here, the RDAP extension for DNS TTL values. Once again, stuff that's gone into an EPP service now coming out in RDAP. And this one is we recently did an EPP extension for DNS TTL, time to live, how long your DNS records to be allowed you to change the TTL of DNS stuff that you put in via EPP. And this allows you to discover that information that has been put in via EPP. It allows it to be discoverable via RDAP. And so, this allows that to come out by RDAP.

This next one is related to the extension registry for EPP extensions, the IANA extension registry for EPP extension. It says it's past, it was at Working Group last call. An issue was discovered in Working Group last call. That's being worked right now. And so, it's probably going to get updated. This will be a BCP document. There's been for a very, very long time, an EPP extension, an IANA registry for EPP extensions. We've all been using it.

This is how standardized extensions for EPP are published and known and managed. This BCP will clarify how those things get put in there, how that all happens. There was an issue raised in Working Group last call that's currently being discussed in the list. And it's been kind of going back and forth. I don't quite know if it's resolved or if it's going to be discussed in Shenzhen.

Next slide, please. So, what's on the agenda in Shenzhen? So, the first one is another RDAP extension. This is also in the RIR space. This is to be able to pull out and get RPKI data. This is if you're using RPKI on your routing, so you're doing route author, route origin authorization, or if you're doing AS provider authorization, aka ASPA, or you're doing X509 RPKI profiles, and you've got that stuff provisioned in the RIR, this allows you to pull it out via RDAP, right, to be able to query and discover all that stuff for the, it's important for the RIRs, as well as other local internet registries and such. I believe that one is also Yazdeep's.

Then the next one is, and H is a formal documented way to do EPP over HTTPS. Most folks today are operating EPP over TLS. That's the way that it's been done since early 2000, late 90s, early 2000s, back when dinosaurs roamed the earth, and we had just tamed fire, and Scott Hollenbeck just wrote it. But now this is a way to do EPP transfer over HTTPS.

This is very important if you want to do cloud deployments for EPP, because it's much easier to get HTTP network connections from your service providers. Context here is that there's another Working Group going called RPP, which is, RPP? RRP? It's my network. What's my outlook? It is RPP. Yeah, sorry. So, many acronyms. Which is about doing an HTTP-based, completely different protocol than EPP, where this gives us HTTP, just EPP over HTTPS, where you can run the same good old EPP, but only do it over an HTTPS transport.

Next slide, please. Next one is about, once again, digging stuff out of RDAP. You're detecting a theme here, right? This is like the wave of, you know, now you can tell that RDAP is common stock in trade. Now people are extending RDAP and digging more stuff out.

This is a document from our good friend Mario Laffredo from the Italian registry. And here, they're doing work, if you've got verified contacts, so you've done some degree of my words, but not necessarily Mario's words, KYC, Know Your Customer about a contact, and how can you get more detailed information about a contact, about inclusion of the verification status information in the contact fields for things such as email address and phone numbers, so you have better idea about the data quality.

This is probably less important to most gTLDs, but probably quite important, excuse me, to certain ccTLDs. It may be important for certain gTLDs that operate in certain special circumstances. I haven't talked to Craig about this, but perhaps like bank and insurance might want to do this. I have no idea, but it's more important for ccTLDs.

The next one is about Extended Key Usage, AKA EKU. I've been waiting all morning to say that. Yeah, AKA EKU. There's been some changes going on. Most people probably don't track this, but there's some changes going on that the certificate authorities are making changes to client certificates that are removing something called extended key usage for most of the certs they issue.

That matters because the client certs won't be usable in certain contexts around this, and so this document which Zayd is working on clarifies and documents a way to do mutual TLS for EPP in the context of these new client certs, so this is actually worth taking a look at for registry operators and for registrars.

Next one, EPP mapping for DNS time to live. This is a document that -- DNS time to live, we talked about it previously. It's a relatively recent RFC 9803. Gavin put this together, and it's a good work around allowing people to manipulate the TTL for DNS records.

Interestingly, someone at the University of Virginia found a minor glitch in the schema, one of these regrettable things where in the EPP schema, they were doing some work, and they found a glitch in the cardinality of one of the parameters, and Gavin, and I don't think that Gavin's online, but Gavin sort of got the black squiggle of frustration over his head, and Jim Gould was like, bless you, Gordon, Jim Gould was like, yeah, you're right, and Gavin was like, ah, dang it, and so we're going to issue an errata on that, and it's going to get fixed, and it won't make really any difference to anybody implementing it, but if you're doing a conformance suite and a test, it makes a big difference. Okay, no big deal.

And then last but certainly not least, and very timely, this got mentioned earlier by Michael, there's a draft getting discussed which Jim and Michael are collaborating on about same entity support for EPP related to all the IDN stuff. It's a delightful read,

and you can see just what it takes to specify in EPP language all of the stuff that Michael was talking about, and so that's also getting discussed, and I don't know if Michael is going or if he's going to be presenting remotely, he's not going, so maybe Jim will be presenting that, and Michael and I will be online watching and asking questions and stuff

So, that's kind of what's on deck for IETF and Regex. As always, a reminder, anybody and everybody is welcome in the IETF to attend the meetings, whether in person or remotely, you've got to pay some money, but to participate on the mailing list is free as in beer, as the saying goes, we don't serve beer at the meetings or anything, but to participate on the mailing list is free, and if you don't know how to sign up, you can ask the Google or me or Michael or Jim Galvin, and so we always welcome participation and engagement from both registries and registrars, and it's a good group to get involved in and learn a lot and also broaden your base of knowledge and history and context. That's it. Happy to take any questions.

ROGER CARNEY

Thanks, Rick, and just to add on to that, and the work at the IETF, they do have meetings, and as Rick mentioned, you have to pay to go, but no work is confirmed anywhere except on list. So, if you're paying attention to the list, all work is confirmed on there, so you always have the opportunity to at least put your input in on list, and

it'll be looked at, so any questions for Rick? Rick's going to ask a question to himself.

RICK WILHELM: One quick thing. Jim Galvin kindly did the slides. He did not give me any notes, so any errors are my own, just to clarify that, but many thanks to Jim. Jim's one of the co-chairs of Regex, and thanks to Jim for doing the slides, but all the notes and errors are my own. Thanks.

ROGER CARNEY: Thanks for that, Rick, and thanks for providing the update that Jim does provide to Tech Ops, so again, please come to the Tech Ops meeting because Jim will do this every couple months or so for us, and he does this exact same makeup for the Tech Ops group. Any questions? And as I noted in chat, it's next Tuesday is the Regex meeting at 16:00 local time, 8:00 UTC. Okay. I think that was our last thing on our agenda wrap-up. Okay, maybe I'll turn this over to Marc to see any thoughts from him on this before we close up.

MARC ANDERSON: Thanks, Roger. Can you hear me okay?

ROGER CARNEY: Sound good.

MARC ANDERSON

All right, great. First, a big thank you to our presenters. Michael, Rick, you both did an excellent job. Thank you. I know there's a fair bit of work and preparation that goes into these, so I think we all appreciated that, and we all benefit from your work on this.

I want to give a little bit of an ask for people to think about what you want in future TechOps meetings. As always, our meetings are only as useful and important as everybody in the group makes them. So, if you have ideas for what you want to see at future meetings or anything that you would like to present, please reach out to Roger and myself. We are always looking for new ideas and new presenters. Other than that, I don't think I have anything to add. Roger, back to you.

ROGER CARNEY

Great. Thanks, Marc. Yeah, and also, thank you to Peter during our first session for providing that and to everyone that contributed. I know Jim does put some time into his updates, and it's great for us to see, and we benefit for his time on that, so that's great.

I don't have anything else. If no one else has anything, we can wrap a few minutes early for lunch. Okay, great. Well, thanks, everybody, and thanks, everybody online, and we will talk to you soon.

[END OF TRANSCRIPTION]