

IDN Variant TLD P2.1 Meeting Minutes 22 August 2012

Meeting Attendees:

LGR Volunteer Team 22-Aug-12

Akshat	Joshi	Present
Alexei	Sozonov	Present
Alireza	Saleh	
Chris	Dillon	Present
Daniel	Kalchev	Apologies
Dmitry	Kohmanyuk	
Fahd A.	Batayneh	Apologies
James	Seng	Present
Joseph	Yee	Present
Linlin	Zhou	Present
Mahesh	Kulkarni	
Mirjana	Tasic	Present
Nadya	Morozova	Present
Neha	Gupta	Present
Panagiotis	Papaspiliopoulos	Apologies
Sarmad	Hussain	Present
Syed Iftikhar H.	Shah	
Vaggelis	Segredakis	Present
Vladimir	Shadrinov	Present
Will	Shorter	
Yoav	Keren	Present
Yoshiro	Yoneya	Present
Zhoucai (Joe)	Zhang	Present

Team Observers:

Edmon	Chung	Present
Raymond (Doc)	Doctor	Present
Rinalia Abdul	Rahim	

Consultants:

Michael	Everson	Present
Asmus	Freytag	Present
Andrew	Sullivan	Present

ICANN Staff:

Francisco	Arias	Present
Dennis	Chang	Present
Kim	Davies	
Dennis	Jennings	Present
Patrick	Jones	Present
Xiaodong	Lee	
Karen	Lentz	
Nicoleta M	Munteanu	Present
Naela	Sarras	Present
Steve	Sheng	

Dennis Jennings: Welcomed the team for the first call. Purpose of the call is an introductory one to begin to set the scene for the major meeting, next face-to-face, next meeting next week in LA. We also have the Adobe Connect, which I encourage you to join and we'll try and use that facility to signal when people want to talk and so on.

In terms of the face-to-face meeting, I think all the information has been circulated on the email list. If you have any questions, please raise those on the email list so that everybody's clear for next week.

Naela Sarras: I can see here from the attendees on this call, thank you everyone, that a number of you have been involved with this program. So I'll try to be brief. This project started out in 2011, early 2011 where the board resolution that we needed to do more on trying to define variant issues so that we may lead to directing variant solutions to the -- managing variant solutions. And so the project started with six case studies that we conducted in different scripts. I believe they were Arabic, Chinese, Latin, Greek, Cyrillic, Arabic, and Latin.

: So there were six case studies that we conducted with the community, much like this format that we're going to use for this project, where we had community volunteers that led the team. And for a period of about five or six months they studied issues related to variants in their scripts. The output of that work was an issues report that was published for each of those six. And that happened around September, October of last year, 2011.

Then we took the combination of those six scripts, put together yet another volunteer team from the community, and that team serves as advisors to the ICANN staff and an integrated issues report was put out based on those exchanges and the information from the six groups. The integrated issues report attempted to synthesize and find what's common and what's unique to each of the scripts that were studied. And that became the integrated issues report that was released in December of last year.

In that issues report, we also had a number of issues or a number of -- basically it talked about what still needs to be addressed before we head into solution area. And so this produced what we're working on now, several projects that materialized out of the integrated issues report. One of them is what we're doing here, which is creating a tool or a process for defining labels that (inaudible) defined as (inaudible). And that's the project we're working on here.

There's a couple other projects that are going on as well. One is examine further the user exchange for activating variants in the root zone. And another project for creating a tool in which variants can be something more uniform that everyone can use, because right

now we have several formats in which the IDN table, for the lack of a better word right now, and the variants are defined.

But where we're at now is the result of all of last year's work, ending with integrated issues reports, identifying the further work that we need to do. I think that captures it, Dennis.

Francisco Arias:

The goal of this project is to develop a process to be used in a later phase. A process to construct and maintain the level generation rules for the root zone. The idea of this project is to build a common framework that will be shared and on the different trip or language or other base communities for the group.

In this project, the idea is to define things like, for example, what the communities should be. Should there be script communities, language communities, or something else? And a few other parameters were mentioned in the integrated issues report. Like, for example, how comprehensive should be the (inaudible) of the code points such have been proved in the LGR. And where there should be some expert view and so on. The idea is to have the group of consultants that would be leading the way and particularly understood along with the main writer, together with Asmus Freytag and Michael Everson.

They would be putting material in front of the volunteers and the rest of the team so we can review and propose changes as we see fit. We also have some plans in (inaudible) review were and they (inaudible) in things more advanced in order to -- for him to provide a specific point of view.

And as you may know, we're planning to have a meeting next week in LA and another meeting in Toronto, just before the ICANN meeting, to call this to produce the final process by the end of February next year. So we have a few months and quite a piece of work to do.

Andrew Sullivan:

I've circulated a couple of drafts, the most recent of which has a file name Proc4LGR20120817A. So that's the one that we're going to talk about today. We have another update in [train], but I haven't circulated it yet.

The basic idea here is that we're going to write a document that is the procedure by which the roots label generation rules are developed. And you may recall from the integrated issues report that we define the label generation rule as having basically two components. One is the repertoire of allocated Unicode code points that are permitted in new labels for the zones. And the second component are the code point substitution rules that would be applied to any of those code points.

So if the code point appears, then there's a set of rules, which may be an empty set, that generates other labels and then those become the variant labels that are either allocated and may be activated, or else that block other labels in the zone. And this particular one is dedicated only for the root zone, so we don't have -- we're not trying to develop something that we would apply to everything. I'm sorry?

So, the goal of this is to produce a procedure by which those rules are established for the root zone. So these are essentially the ground rules.

The draft that we have depends critically on two other documents that if you haven't read them, you're just going to be completely at sea. So I encourage you all to read them. The first is the integrated issues report from last year and there is a reference to that in the notes in the references for the -- in the draft. And the second is an Internet draft that comes -- the Internet Architecture Board has adopted called "Principles for Unicode Code Point Inclusion in Labels in the DNS." And that's available from the IETF tools site and I believe I circulated the references of those to the mailing list as well. So, if you haven't read them, you'll need to.

The integrated issues report is lengthy, but there are a number of terms that we're going to depend on from that document and the background is in there, so I'm not going to repeat all of that in this document, because we've already published it. The IAB draft is quite short and I think readable, so it should not take a lot of time to read.

The IAB draft includes a number of principles, which have been summarized somewhat - a little bit loosely in this draft. It's those principles that appear in section 3 of the draft. I think that those principles are important foundational ones, but I want to point out that the IAB has not published this draft. It's not an RFC yet. And so it has not achieved even IAB consensus, never mind -- well, it's never going to achieve IETF consensus because it's an IAB draft. But it hasn't been out for wide comment and so on, so that draft could change, but I've taken those principles as foundational anyway, so if people want to comment on them, that would be a good thing to hear if people don't think that they're a good idea.

I'm not going to go through the draft in detail, but I just -- I want to point out sort of the basic structure of what is composed in this. The basic structure is that we take those principles as foundational principles and notice they're principles, they're not rules strictly, but they're principles that are supposed to guide individuals in making decisions.

And then we have a sort of two-phase project. And the first phase is groups of people who are interested in or otherwise engaged with a particular language community, a given script or a writing system. So this is a sort of linguistic community very loosely defined. It's a little amorphous because I think one of the things the issues report highlighted very, very clearly last year is that different kinds of systems demand difference sorts of attention. So the kinds of things that you can do for something like [han] with a very, very limited number of languages that use it, very different from the sorts of things that you would want to do for say, Latin, which is like half the languages on earth.

So we have a very different kind of response for each one of these things and I think that that's appropriate. But the idea is that you have different bodies that work on these things and those are subsections of the entire Unicode repertoire.

Then you have a second panel. And that second panel is ultimately responsible for producing a single unified set of rules for the root zone. So that second panel has to filter essentially all of the proposed inclusions for the repertoire and all of the proposed rules. And that job is one that is taken by that panel for the entire root zone. So what it produces is the set of rules finally for the root zone.

You may remember from the integrated issues report one of the key things that we emphasized in that report is that because there is one zone there can only be one set of rules. That's true by definition. And that means that the set of rules needs to be consistent regardless of the linguistic community for which the rules are targeted and that means that you have to take into consideration things that are perhaps not true for a particular group of language speakers or a particular script. There might be cross script effects and so on. So that's the job of the second panel.

And then what we want to do in the final two pieces of this are to evaluate this proposal against the principles to see whether this proposal actually can conform to all of those principles. And we want to sort of describe in terms of those parameters that we described. In the integrated issues report, we had three parameters, but Asmus in particular has been very helpful in describing a fourth parameter, which is this matter of centralization. And so, we want to describe it in those terms and sort of understand what sorts of settings we are establishing.

So I think that's sort of an overview on -- I don't know if people have comments on these things. Oh, I should also respond to some questions that went to the list this morning from Joe Zhang, who asked what the outcome of the project is. The outcome of the project is this document, as I think Francisco already said. Dennis, of course, is the convener. There aren't exactly going to be any subgroups. There's only one document that we're working on here. I'm going to be the editor and the fourth question that was there was when the primary panel should be established and so on. That's all stuff that comes after this. So only after this document is adopted and accepted by ICANN and by the board, I guess -- I guess it's the board who does it. Somebody from ICANN would have to tell you that. Ultimately the procedures that are outlined here would be implemented later. And that's sometime early next year I guess is what the project plan calls for.

Alexi Sozonov: I've read the document and it's very well structured and the ideas it's promoting is very reasonable ideas. Do you think it would be a good idea if like each of us after reading the document we should just in the email at least have just our thoughts about it, that these were like broadly approved or just we like the idea or we have some suggestions?

Dennis Jennings: Alexi, yes, I would encourage you, (A) to read the document and (B) to raise any questions or comments that you have on the list. The more work we can do on a continuing basis on the list, particularly over the next couple of days in preparation for the face-to-face meeting the better. So Alexi, please read it, if you haven't already read it, raise your comments, submit your comments to the list, so that everybody can see them, and your questions. Thank you.

Joseph Yee: More to Andrew for these questions. I guess there's actually just to procedure, but the document mentions that the evaluations were trying to do across the script evaluations. And it seems that to be -- you don't need to be cross script with every language just crossed over each other. But mainly just to cross script the particular language/scripts to nothing, almost to the ASCII point of view. Is that a (inaudible)? Or will you be (inaudible) or you mentioned to really like (inaudible) of every language against -- every script against other scripts.

Andrew Sullivan: Fundamentally there is -- when you accept an allocated code point for use in new labels in any zone, what you are saying is that that thing is an acceptable code point for your use in that zone.

In the root zone, we have a problem in that normally what you want to do is accept things that are sort of narrowly tailored to your user community. And the root zone -- the user community is everybody on the Internet. So in effect, what you have to do is when you are accepting a code point for use in the root zone, you have to evaluate that against all of the other code points that could possibly be used.

Now, obviously if what you're looking at is, for instance, a [han] character, you don't actually have to worry about how that's going to interact with, say Latin small letter A, because they're just completely different uses, right? There's not anything that -- nobody uses those sort of together in a string or something like that. They're just not -- they're not portions of Unicode that collide with one another. Although of course, in modern Internet usage, the Latin portion that corresponds to ASCII is in wide use and sort of needs to interact somewhat more.

So, we expect that every panel, including the secondary panel, is going to have enough experience with all of Unicode to understand how the different pieces fit together, at least in the abstract. So this isn't a sort of combinatorial thing where you go through everything in Unicode and then evaluate it against every other character. That process would never complete. But you do have to evaluate for each thing that you're going to permit, you have to evaluate whether it's going to have any effect on any other code point that you're going to permit. And you have to take into consideration of course the things that are already in the root zone as a result.

So I hope that answers your question. It's not -- there isn't a sort of either/or answer here is I think what I'm trying to say.

Chris Dillon: It's actually just following up on what Joseph was saying. I mean the question that I had in my mind as I was reading it was, I mean may be very difficult to answer, but it's whether the similarity is -- to what extent we're thinking of similarity as we're creating rules? Or whether basically what will happen is that rules will suggest things for the root and at that stage you have another process, which comes in and analyzes now to what extent have we got similarity. There are two possible times when similarity tracks can be (inaudible).

Andrew Sullivan: Right. So, the idea here is that the string similarity -- this is a consequence of the decision we made in the integrated issues report. In the integrated issues report, we decided to say that string similarity, like the full string similarity and all the rest of it, is just going to remain the responsibility of the string similarity panel. So it kind of gets punted on. And what this process does is it establishes how you do code point substitution rules only.

Now, there is still a similarity issue there, right? Because if you look at say, Latin A and Cyrillic A, I mean you've really got a serious similarity problem there.

Chris Dillon: So, yes, they're identical.

Andrew Sullivan: And at least in just about every response that we can think of. I'm sure you could find a way to distinguish them, but that's not going to be something that you can distinguish code point by code point.

Michael Everson: You cannot find a way to distinguish them. For the record, there is no way of distinguishing them. No one ever, had ever done that, ever. And no one ever will.

Andrew Sullivan: Right. So, I would -- it seems to me that the problem there is going to be you're going to need some kind of automatic rule for those things. And whether it's an automatic rule about context and other things in the string, or whether it's going to be a straightforward substitution rule is a live question.

The draft -- I think this is in the draft that I circulated, actually says you're not allowed to have context rules out of this process. And one of the consequences of that, and I urge people to read this carefully because it's -- this is going to be important. One of the consequences of that is that things like substitution -- or things like derivation of substitution of Latin A and Cyrillic A might be something that the secondary panel is going to insist on, precisely because it can't say well, if you have a Cyrillic A, you have to have at least one character in the string that is clearly not confusable with some sort of Latin character. That's the sort of context rule that would be not allowed, would not be permitted under the text as it's currently written.

The reason for that is one of the principles, the simplicity principle I think it's called. I forget the names of these, even though I participated in their writing. Yes, the simplicity principle says that really complex rules are to be avoided because it makes them hard to implement. And so that means that very complicated context rules are not going to be allowed and that would mean that you can't do sort of whole string evaluations in the rule set.

And that means that we've got to pass this on to another panel and then they have a problem. So that's one of the areas where I think a lot of consideration needs to be put. I mean this is a proposal and I'm not wedded to it, but I'm trying to conform to these principles that the IAB seems to like.

- Chris Dillon: Okay, so the summary of that is that the code point similarity is in the rules and the string similarity is (inaudible) off to -- it's in a separate process.
- Andrew Sullivan: I think that's right.
- Akshat Joshi: Yes, regarding this substitution principle line, when we said that a four point substitution would be considered [near] end string similarity would be routed to some other panel. When we see a string, do we say a [plethora] of character might (inaudible) with a plethora of character? Or we are essentially talking about the actual string at the period (inaudible) applied.
- Andrew Sullivan: Because the integrated issues report sort of kicked string similarity off to the string similarity panel, which is already an existing process within ICANN, I haven't attempted to -- I really haven't attempted to think about what their rules are, because it seems to me that they could make whatever rules they want.
- So there are two possibilities that I can think of. One is that they could have a rule that the string that they are considering must not be similar to any actually allocated string that's already been allocated. And by string, I guess what I'm talking about is a candidate U label, right? So we're not talking strings sort of absolutely, but rather things that could actually be U labels because there are lots of strings that are confusingly similar or whatever. Like for instance, prior to normalization. And none of those count because they're not even candidate U labels and so we're just not going to worry about them.
- The other possibility is that they're going to think about strings -- all the possible strings that could be confusable. Those are two things that I think are outside of the remit of this group because we're only talking code point by code point. I think that it's an important consideration and I think it's something that we might want to sort of have a secondary note or something like that to point out that there's this entire set of issues that was -- that we've considered to be outside of our remit and therefore we haven't talked about them. But probably the string similarity guys need to come up with a set of answers about this because really you're only going to be able to evaluate the safety of the entire root policy when you take into consideration both of these portions.
- Akshat Joshi: Just one clarification, Andrew. I would like to have in the point number 3 point, when we are (inaudible) the procedure and for the things that primary panel is supposed to do, is that (inaudible) two categories of core point execution. [One is (inaudible), that is without any context. And the second one is one is too many or mainly a two one, without any context.] And again, my -- well, (inaudible) that portable (inaudible) where we have a core point, like many are supposed to be subject to (inaudible). So in that case, it would not be a case of candidate you level asset, but it is (inaudible) character similar to (inaudible) characters.
- Andrew Sullivan: I see. So what you want is a series of characters, not really a string, but a sort of sub string that then corresponds to a different set of things. And you can't do it code point by code point.
- Akshat Joshi: If we have one single core (inaudible) mapping to a cluster, but it can -- it will be a sub string matching to a sub string.
- Andrew Sullivan: Okay. If you could propose text for this, it would be helpful. And I think we should do that on the list. But if you could suggest a way to write this compactly, it would be helpful. The key rule here of course is that you want to make sure that this doesn't violate the simplicity principle as well. So that it needs to be simple enough that somebody can understand it. So let's work on some text for this and we can do that on the mailing list. How's that sound?

Akshat Joshi: Sure, I'll do that. And just one thing like -- actually I'm talking about there will be kind of (inaudible) that we are mentioning that there were (inaudible) issues to report. And as far as our perception goes, it's not much a complicated thing to understand for the native user of the language.

Andrew Sullivan: Ah, but wait a minute. You've just raised a really important point. It doesn't matter if it's complicated for the native user of the language. The key thing is that it needs to be not complicated, even for people who don't speak the language. That's part of the rule of this -- that's part of the simplicity principle here. That in the root zone in particular, we cannot depend on people understanding the rule because they already speak that language. In that case, we're never going to get anywhere, because for any language you pick, somebody doesn't speak it and we're talking about the entire world here, right?

So the rules need to be simple enough that if you understand enough about Unicode, you can understand the rule, even if you don't speak the language.

Akshat Joshi: Yes, actually I would like to rephrase (inaudible). It would be like a native user of the script actually. So in that case, it becomes independent of the language in this case -- in this particular case.

Andrew Sullivan: Even the script doesn't matter. You need to be able to understand this if you understand how Unicode works. And that is a very high bar. I am fully aware of how high a bar that principle is. But there's a reason that rule is there or that principle is there. And that principle is there because we're talking about the entire world. This needs to be -- the ultimate set of rules needs to be something that can be maintained by people who don't speak the language and who don't use the script and so on. It needs to be something that can be understandable to somebody who understands enough about how Unicode works and who understands enough about how DNS and IDNA work so they can understand what the rule would be, even though they don't understand why the rule is there.

Akshat Joshi: I think we better take it on the list so maybe we'll (inaudible).

Andrew Sullivan: Yes, but I mean you're leaning on a very, very critical point, so I mean I'm glad that you raised it because it's one of the things that I'm most nervous about in this. People need to understand that those principles are very, very conservative and they are more likely to restrict code points than they are to allow them.

Akshat Joshi: Sure.

Raymond Doctor: Yes. I'm coming back to the simplicity principle. And if you read it very carefully, the first sentence itself says easily understood by users with only some background. Now what does this mean? Background of language or background of Unicode? I think Andrew means more background of Unicode than background of language. So I think somewhere some amplification is required because to me the simplicity principle seems rather weak enunciated right here.

Andrew Sullivan: So, two things about that. The first thing that I would say is these principles are sort of short descriptions to remind people and if you want to read the full text, you need to go and read the original Internet draft. That's something that I try to state as clearly as possible in the text, but in case it isn't clear, I want to emphasize it again. I'm going to post in the chat room the -- there. In the chat room I've just put a link to the original document. And the principles are rather more extensively discussed in that document. And I think that they're sort of the thing that you really need to read because these two sentence summaries are here to remind people what the point of these principles are, but the real document -- I'm not going to reproduce all the text there, because that's the entire document. That's the entire Internet draft. But I think that that's a key thing.

I agree with you however, that only some background is probably a little bit too big. So I'll see if I can add something to clarify that.

- Raymond Doctor: Well, but Andrew, what I'm referring to is the basic IAB CP document itself. I went through it and the simplicity principle still, for me at least, doesn't make much sense. It still remains rather vague. I think it needs a bit of amplification and clarity. What you stated here is in essence what the principle says actually. But it isn't a very clear illustration of what should be and what should not be. This is what I've been trying to say. I think this is what Akshat is also trying to stress.
- Dennis Jennings: One is that that's feedback to the IAB. And two it sounds like a question that we need to explore in more detail next week at the face-to-face meeting. Andrew.
- Andrew Sullivan: Probably true, but I want to be very careful. These principles are never going to give us an algorithm. They're never going to give us a way by which you can put stuff in and get a guaranteed output where everyone's going to agree. That's why they're principles and not rules. If they were rules, we could just plug them in and sort of turn the crank and we'd get the answer. But they're not an algorithmic set of rules, they're just guidance. And so to a certain extent, they're always going to be vague, unfortunately. And there's a tension actually among some of these principles that is something I don't think can be resolved very easily. So I appreciate the point, but I think that we're going to have to be prepared for a certain amount of vagueness in these principles because they're intended as guidance and not as rules.
- Raymond Doctor: Yes, because that's just what I wanted to say. Some of the principles conflict with one another. It's something like the Gricean maxims. If there's a conflict, then somewhere you do establish a hierarchy and say this one is more dominant and this one is at (inaudible) most of the year to the other one. Like a super maxim and a maxim in Grice.
- Andrew Sullivan: Yes, but ultimately a certain number of these things are just going to compound the judgment.
- Raymond Doctor: Yes, I understand.
- Andrew Sullivan: If there is a primary principle in all of this, I think this is clear also in the IAB document and certainly I've tried to make it clear in this document, if there is an overarching principle, it is a conservatism principle.
- Raymond Doctor: Absolutely.
- Andrew Sullivan: And the reason for that is simple, right? It's always easy to permit another one later if experience tells you it's possible, but it's just about impossible to take something out once you've allowed it.
- Edmon Chung: One, ultimately the issue of whether contextual rules are allowed, I think I heard a couple of people have some feedback, I guess initial feedback that they might need to be considered. But I haven't really heard, besides the one particular example, why we really need (inaudible). So here's one of them that I think perhaps needs to be considered. What you mean by simplicity is that -- I mean how does that compare with the bar if we say that there are certain rules where the first character and last character need to be or cannot be certain characters. Because I believe some of the language policies do have that type of characteristic. Would that be permissible?
- Andrew Sullivan: That is a context rule, right? The rule there is essentially having to do with the position of the character with respect to the rest of the U label. And under the current proposal, the answer is no, that's not allowed. That's some other problem outside of these things.

Now, one of the difficulties that I see with that thing is that the basic idea in some cases is that you're going to say this character in a final position, for instance, is blocking or is going to be blocked. And so that entire label would then no longer be allowed. And one of the difficulties with the proposal as it currently stands is that it doesn't allow for that kind of flexibility, essentially allocation is all the way and we don't have sort of easy rules about required to activate versus not required to activate and so on.

Maybe the current proposal errs too far on the side of simplicity, but I really want us to think very hard about the consequences for implementation of very complicated rules when we have to cover the entire world. It's quite a different matter in a zone where you know you are sort of dedicated to some linguistic group or some geographic area or something like that. It's a very different matter from the case where you have to deal with everybody in the whole world and have all of the rules that would cover everything that you could possibly do in Unicode. And the position that the current draft takes is that because this is having to do with everybody in the entire world, therefore a bunch of things that would be normal and permissible in a particular zone are things that are just not going to be allowed in this zone. And that means that a large amount of flexibility and a large number of cases that we would normally want to permit for any -- for this or that language are things that we're just not going to allow because it makes the rules too complicated.

So what this really is, is everybody has to give up something in order to have a common root. And because you have a common root by definition, it's just part of the consequence.

Edmon Chung: I think we're getting onto a very important and interesting point and that's why I want to intervene again. When we talk about this simplicity principle, we also need to think about the context as well. In the DNS context, we always -- I would say there's always a rule, certainly agreed, that -- and the person's character has certain special rules, like it cannot be the hyphen for example. And when we talk -- we go down a root of that -- even this particular -- this type of contextual rules cannot be set, then we might need to go back and look at the languages that need those kind of rules and we might even have to take out those characters that might have that characteristic, if they are put in the first or last place, and take them out entirely of the permissible (inaudible), which is what we did for hyphens. Actually we didn't. Hyphen is the last (inaudible) in the TLD, right? Is it?

Dennis Jennings: Edmon, I'm sorry to cut across you, but I think you're raising a crucial point here.

Edmon Chung: Yes, it's pretty important, so.

Dennis Jennings: Yes, so I think that we can put that on the list of things to begin to discuss when in LA.

Asmus Freytag: The one important thing is -- that's an issue for the next version of the draft, we need to go from just pure principle, and the abstract issue is a high bar, to looking at the specific restrictions that we think follow from these principles and evaluating them one by one. We need to evaluate the context rule limitations in more detail. We need to explain what kind of context rules are too complex, et cetera. I think just saying it's not simple is, I think, too simplistic a way to work in this style.

Dennis Jennings: So as a general thing, Asmus, you're saying that we need to do what in physics we used to call the Gedanken experiment. We need to take cases and see how they stand up under the rules that we're proposing.

Asmus Freytag: That would be one good way of doing that. It would be very specific in terms of showing how the principles apply. And that's the only meaningful way how we can, as a collective, get to a decision of what we want to propose as a process.

Naela Sarras: Nicoletta sent an email yesterday, our time, detailing all the things that the travelers need to know coming into the meeting. We tried to put the meeting schedule in there, but also the logistics of where everyone is staying and how to get to the new offices. And again, We've moved to a whole new office. So take note of our new address and make sure your taxi goes to the new address.

If there are any questions, Nicoletta and I both provided our mobile numbers. Please feel free to call us anytime when you get in and we'll do our best to help. But take note of all the details sent out last night. If anything is missing, email us and we'll try to answer the question as soon as possible.

Dennis Jennings: Thank you very much indeed, Naela. I raised a question on the list about what if it's 3:00 a.m. in the morning and I need a bit of help? I don't particularly want to call you or Nicole or Nicoletta in the middle of the night. Is there an emergency number or is there a -- or do I just call you? What's the story?

Naela Sarras: Just call us. If it's an emergency, just call us, please. It's okay. We have security staff that we also call into emergencies and we have access to them. So call us and we'll make sure we'll get something.

Dennis Jennings: If you would, type into the mailing list, your question and we'll pick it up on the mailing list. We very much appreciate it and I look forward to seeing most of you and hearing all of you next week at the face-to-face meeting in Los Angeles. Thank you very much and good-bye.