ICANN
Transcription ICANN64 Kobe
GNSO – NCSG Open Meeting Part 2
Tuesday, 12 March 2019 at 17:00 JST

Note: Although the transcription is largely accurate, in some cases it is incomplete or inaccurate due to inaudible passages or transcription errors. It is posted as an aid to understanding the proceedings at the meeting, but should not be treated as an authoritative record.
The recordings and transcriptions of the calls are posted on the GNSO Master Calendar page
http://gnso.icann.org/en/group-activities/calendar

Raoul Plommer: For the record, Raoul Plommer. I think that was a great presentation and myself am now more interested in the RPMs. I didn't know the situation is as bad as that.

Stephanie Perrin: I'm just so glad you did that, Raoul.

Kathy Kleiman: Is there any other comment, including from members of the Sub Pro?

Stephanie Perrin: And I do -- it's Stephanie Perrin for the record -- I do apologize for nipping out to the bathroom. I missed the bathroom break earlier so I missed some of your presentation. But I was wondering, we had a little bit of discussion with the board this morning about this whole defining the public interest. What do you think are the chances of rolling back the PICs once we start that process? Is there a strategy that we could enact or could work on that would help us kill two birds at the same time?

Kathy Kleiman: I'm going to try the answer and then I'll pass it to Robin to see if she thinks it's the right strategy. What I'm told by people who were there at the time and

who were on the board was that the PICs should never have gone in. These - mandatory public interest commitments by the way are something else. So we're not insulting all public interest commitments, so. But these voluntary ones where you can put anything you wanted in, there are some board members who think they won't be part of the next model registry agreement, the agreement everyone should sign, and that they are outside the scope of the new bylaws, which have come since.

On the other hand, what's going on in the Subsequent Procedures working group dominated by the incumbents is oh we had voluntary PICs. We're going to grandfather them in. We're going to continue them and yes we're going to make them all enforceable by ICANN, which makes ICANN then the content regulator. It's - and I - it's really hard to know how to stop them boulder.

Stephanie Perrin: And if I can just to a follow-up question to that. Did we get anything out of the competition and consumer trust review that would - because to me these are anti-competitive, you know? There's no - it's pretty clear in my head and that we - ICANN has a very strong responsibility for policing anti-competitive behavior in this area and so that review committee, which I didn't have enough bandwidth to follow, should have looked at that issue.

Kathy Kleiman: I think we spent a lot of time knocking them out of content. I mean they like content for consumer trust, but I'll go back and take a look. It's a good question. And, Robin, I want to know your thoughts.

Stephanie Perrin: We had somebody on that. Who was it? On the Consumer and Competition Review committee? Carlos, yes. I thought so.

Robin Gross: Yes, I just wanted to say something about these PICs that Kathy's brought up. I'm really concerned about them as well and one of the reasons is because when we, like she mentioned, we sort of grandfathered them in all when we did the whole new bylaws. And it defined, you know, ICANN's

mission to be, like, including the things like PICs. So I'm - I mean at the time I was like, "Oh, my God. We're locked in forever and ever. How are we going to get around this?" So it's a big problem and I don't - I'm not sure how we're going to get around it without, you know, saying you have to change the bylaws in some way or something.

Collin Kurre: This is Collin Kurre. So I do think that there's a big problem with public interest in general in ICANN and the striation of PICs is also a problem in and of itself because I haven't been - I wasn't around for these kinds of conversations that happened years ago but just trying to do research into voluntary or mandatory PICs really highlighted the problem there that conflicting definitions or interpretations or the existence of different levels of public interest commitment was very confusing, which makes - which renders the term rather moot in my opinion and it makes really difficult for people who have legitimate objections to make use of the limited public interest objection, for example.

So I am of the opinion that a revision to at least Spec 11 is absolutely mandatory in order to bring it in line with, for example, the forthcoming human rights core value and I also think that there should - public interest commitments could be consolidated to ensure that they can be properly monitored and enforced.

Robin Gross: Can I just have another go at this public interest commitment issue because what we're seeing is, you know, policy is supposed to be made from the GNSO through PDPs and from, you know, the GNSO process. So that's how policy gets made but then you see the GAC and ALAC lobby the board for - to change what the GNSO came up and so they do. They just say, the board just says, "Okay, well it's going to be a public interest commitment then."

And then it's totally, you know, reworking what came up through the GNSO process based upon the GAC and the IPC lobbying the board to get the things that they want thrown in at the last minute under the label of a public

interest commitment. So I share this concern that it's not really public interest that they're doing, it's just more about, you know, trying to appease powerful actors by giving them what they want so they don't cause problems for ICANN.

Kathy Kleiman: I think it's David and then (unintelligible).

David Cake: As, you know, as Kathy well knows I could happily rant for a good half an hour on how much I hate PICs and how it's essentially one of the worst policy both processes and decisions and outcomes of the entire new gTLD program. Like, it's essentially this, you know, the process was essentially the GAC complains, the board sort of makes something up, the outcome was that it literally does the exact opposite of what it's intended in that many of these so-called public interest commitments are literally policy that was rejected as not being in the public interest during GNSO processes and just some, you know, but the IPC really wanted it. They're the public.

So, yes, this is the end at any attempt to make something that is any way sort of broader idea of what is actually the public commitment rather than what the GAC thinks is in, you know, the GAC decides their lobbyists really want is the - any conversation at trying to actually come up with some working definition is just a disaster. Like, we've never had anything that's gone anywhere and we usually have a large number of several people in this group advocating that it would be a disaster to even to get there.

I do kind of disagree on what we really need. We do need to not just - your push for any review because we know it was not. It has never been declared to be consensus policy. Like, no one - even the board has not had the goal to try and claim that it was consensus policy. So of course it should be reviewed. Yes, there's no real argument for it not be reviewed. And so any chance of doing that, we absolutely should.

And I was also just struck by one of the really great ties into earlier discussion about developing really good solid norms is exactly the sort of thing that would help us, you know, find a light that might lead us out of the public interest wilderness to actually have a bunch of norms that ICANN has actually said, "Well, these are - these guide our definition of the public interest."

So I doubt we'll find anyone who's neatly developed a set of norms about, you know, reining in the IPC but in terms of at least I mean the security and sort of things, this would be - that is one of the things where it would be really helpful to sit and go, "Oh we have some actual norms here to give us some idea of what is or is not in the public interest." So it's not just made up on the spot by, you know, whoever decides.

Kathy Kleiman: Just a quick note that we are reviewing this now. It is one of the myriad, the thousand issues in front of Subsequent Procedures, and if we have enough voices we could raise an objection to continuing these voluntary - I've heard them called voluntary commitments, voluntary nonsense and voluntary crap. We can't - but we have to have more voices that we have now in Subsequent Procedures but it is being reviewed right now.

Anriette Esterhuysen: Thanks a lot, Kathy, and this is a really interesting discussion. So just a novice question. So the private content regulation, is that regulation of the domain name as in taking on the domain based on the content that is hosted, which is - that is really extreme. So I would - I mean have you done - is a human rights impact assessment being done because I think that - this - doing a human rights impact assessment of this particular procedure of private content regulation would be I think very easy and actually it's a very doable thing.

You said there isn't really global content. There is. There's human rights law, which has very clear laws about freedom of expression and freedom of association and exceptions and limitations to freedom of expression and

association. And there's a parallel process happing at the moment with the social networking platforms, where norms and procedures are being developed to apply those laws to the content regulation that is being exercised by these platforms, who are in a similar way no longer I think deserving of non-liability because they're no longer just intermediaries.

And so I think for me this also - it touches on that fundamental principle of what is an intermediary and what is somebody who actually is interfering with content to the extent that they could be held liable. And I think ICANN is - this is a really slippery slope for ICANN. ICANN could end up a little bit like Facebook.

So I think there are ways to look at it. Also look at (David Kaye)'s report on content regulation. Last year he was the special rapporteur on freedom of expression and he did a very in-depth report on content regulation in social media and looked at how companies could actually be creating new rules and procedures that comply with international human rights law.

You can also ask (David Kaye) to post a comment on this particular subsequent procedure and he's very responsive and he's very sharp and he's - he would be able to then express, from a freedom of expression perspective, why this is problematic. So yes, this is really - I mean I can just see ICANN really regretting this. If ICANN is not ensuring that the contracted parties are upholding freedom of expression, if that's not actually concretely part of the criteria, ICANN could end up being held liable or being called to account in the way that actually the big social networking platforms are being called to account at the moment.

Collin Kurre: So really briefly, we didn't look at these kinds of, well, garden concerns that Kathy outlined in the demo impact assessment that we did because we were really focused just on these, like, on the procedures, the concrete procedures, but I think that the model that we have could absolutely be applied to this. So if concerned members wanted to give it a test run at trying

to apply this then we could then present the negative impacts and the description's negative impact scenarios and recommendations to (David Kaye) in order to give him the type of background information that he would need to issue a robust and influential statement.

Anriette Esterhuysen: And just a quick -- it's Anriette for the record again -- just it's related but not directly. I have been asking .africa whether I can register downwithpresidentsforlife.africa and downwithdictators.africa. I haven't yet got an answer from them actually and the process is so expensive now. Actually - and I need some money because .africa domains are not cheap.

But when I've spoken to .africa people about their policy they could not guarantee that I would get away with registering those domains. They thought that maybe at some point the African Union, who's part of the whole structure, might be able to block that.

Kathy Kleiman: Thank you for the discussion and to our sub-team - Subsequent Procedure members. We've got our work cut off for us, and let's work together, please, and coordinate. Together we might be able to put up a strong defense.

Stephanie Perrin: Thanks, Kathy. Thank you very, very much. That was a great presentation. I guess we should be hearing very shortly from our guest from the - ah.

Woman: (Unintelligible)

Man: One of them was just here (unintelligible).

Stephanie Perrin: Oh. Wonderful. That's great.

((Crosstalk))

Stephanie Perrin: Well, no. Thank you for coming. It was great. It was great. Yes. Oh wonderful. Please. Have we got chairs here?

((Crosstalk))

Stephanie Perrin: Okay. While we're getting the slides hooked up, thank you so much for coming. My name is Stephanie Perrin and I am the chair of the Non-Commercial Stakeholder's Group. That is the umbrella over the two constituencies and unaffiliated members -- thank you so much -- underneath that umbrella, the NPOC, which is the Not-For-Profit Operational Concerns and NCUC, the Non-Commercial Users Constituency. And so that's who we are here. I won't take up time going all the way around the table I think but if you would like to perhaps introduce your members, that would be wonderful.

Gavin Brown: Yes. I'm absolutely happy to. I should just point out that Ram Mohan is the coordinator of our group for the - he's our unofficial deputy, so he's on way. He'll be here very shortly but I don't see any reason why we can't just kind of quickly introduce ourselves.

So to my right is Jody Kolker from GoDaddy and then starting around the corner there, the gentlemen in the corner is Steve Hollenbeck - sorry, Scott -- and then Andy Newton of (Abba Aaron). Next to him Jorge Conner from (Nicamess) and Tomofumi from (Digicert). And we have a few other members of the team who are scattered around. Not all of them can make it here today. I'm Gavin Brown. I'm from (Central Nick).

I guess I can - I mean I've heard this - I've heard Mom give this talk now five times so I could probably get started, right? Okay. So yes, so we are the Technical Study Group on access to nonpublic registration data and as I said Ram Mohan of (Afflilias) is our coordinator. Let me bring the slide up because I can't remember them.

So we're here today to tell you about what we're doing, who we are and what we're doing and where we are right now in our current thinking. Just to move on to the next slide to ask the question what is the Technical Study Group?

We were obviously created at the behest of Göran Marby to put together a technical model to explore the possibility of how a technical model for access to nonpublic registration data could work.

The motivational background was, as the slide shows, is to try and balance data protection requirements with legitimate interested third parties and try and come up with a way that reduces the liability of contracted parties, mainly registries and registrars, when providing such access.

One thing I want to make clear is that the TSG is a technical study group. We're not a policy group. We're not a - we're not deliberative in terms of policy. We have no opinion on policy. Our approach was to come up with a solution that given a set of assumptions about the policy environment would accommodate a range of operating models and a range of policy outcomes.

So who are the TSG? Well on the next slide you'll see a list of all the people. We've had a couple of face-to-face meetings and our photos from those meetings. I would also like to bring out the great deal of support and help we've had from ICANN staff in our work. It's been instrumental to the whole process and we couldn't have done it without them. Everyone in the group is very grateful for that.

I'd like to just move on to the next slide which is just talk again about some of the assumptions. So when it's important to understand when we talk about assumptions, most - these assumptions are things that we have been given as a priori, things to assume when designing the system. And again we have no position on the validity of those assumptions. We were given them as a basis for which to do our work.

It's a matter for the community to decide based on advice from legal experts on - as to whether these assumptions are indeed true. So I would hope that

you would bear that in mind when thinking about the work we've done that we are technical experts. The other people in the room a group of technical experts and we've bringing technical expertise to bear on a technical issue and not commenting on legal and policy matters.

So this diagram outlines the core assumptions that we've used when designing our model and it's been - the diagram does seem similar to one that's been -- oh here we go -- does seem similar to one that's been displayed in the session on the next steps for the Expedited PDP, the right-hand side of that diagram. And since Steve is not here, perhaps Ram could take over his place and just go through this slide.

Ram Mohan: Thank you, Gavin, and really apologize for being late. I was not cognizant of the fact that I was going to be improving my step count from the other building to here, so really apologize. On the assumptions, the - just to be clear, these are kind of the conditions that we started off with. These were in many cases what we were given. These were, you know, what we were told are the assumptions that we should begin with and starting to implement and/or starting to thinking about a technical solution.

So when we started in our first cut at this I think we came up with seven or eight assumptions but eventually we ended up with 12 assumptions that we've documented. This slide actually shows the - what we believe are some of the core ones. What's listed in parenthesis are assumptions that they refer back to numbers in the document that we've published.

But one of the core assumptions coming in was that ICANN reduces gTLD registrars and registries' GDPR liability. That's - that was kind of an assumption that has been - was provided to us. In addition, there's another assumption that RDAP is the mechanism to get to access to get nonpublic data and that Port 43 is going to be deprecated.

Further, that all access to gTLD nonpublic data will only be via ICANN. You know, that's what we've shown there in that ICANN gateway, but that's another assumption. And that if there queries that come in from unauthenticated sources then they'll be dealt with as per policy and that ICANN oversees the protection of credentials, as well as the validation validity of the credentials.

So these are the assumptions. Again, we did not sit there and say are all of these assumptions true? Are all of these assumptions the assumptions that are accurate? What we did was to say these are the set of assumptions that if we don't state these assumptions it would be very hard for us to actually go make any further progress on thinking about a model.

To take us to the next slide, Andy?

Andy Newton: Yes, so after we discussed our assumptions and wrote them down we took it the next step, which was to create a use case. This is fairly common method for doing systems analysis and design. So we came up with some set of use cases in or user journeys of, you know, who's going to use the system and what nature are they using it.

And so we discussed - we have use cases for users who are authenticated and authorized to have access to nonpublic data but they only need to do it on more of an ad hoc basis. And then we took that one step further and said well what about users who are both authenticated and authorized but they need more than one query every once in a while? And we took that - then we - when looked at, you know, use cases for users who are no authenticated or authorized and what do you do with them?

And then finally the use case of the user that is the actual registrant or the data subject. And then from there we went on to some systems requirements, and some of the requirements are very specific to components of the system and the different components of system basically are around a web portal

that ICANN would run in order to deal with expedited queries or the ad hoc queries that come in and also to help with users with finding their identity providers, if that's a system that is needed via policy.

The other one would be - the other components are an ICANN-RDAP gateway and the RDAP servers of the contracted parties as well in addition to identity providers and authorization determiners. So we have requirements around those. And then we have some system requirements that are kind of broad and global such as best practices for information security, using secure protocols where appropriate, following best practices for controls of data, doing things like logging the queries and then obeying log retention and data retention policies, allowing some sort of transparencies to reconciliation of the logs and overall issues around organizational continuity of the system. That's it. So.

Ram Mohan: Thanks, Andy. Scott, will you take the group through our proposed technical model?

Scott Hollenbeck: Sure. Glad to, Ram. Thank you. Next slide, please.

We did change a little bit, okay. So the model that we settled on is based on three standards-based technologies: RDAP, Registration Data Access Protocol, open ID connect, which helps us address requirements for identification and authentication, and (OOFF) 2.0, which helps us address requirements for authorization.

If you're familiar with the traditional Whois client server model and you look at this picture, you're going to see a couple of other little, you know, server-like entities on the picture and that's because the model has become a little bit more complicated, right? Instead of a client sending queries directly to a data source, you know, as Andy described, we have this notion of an ICANN run and operated RDAP access service.

It provides a proxy-like function and helps manage the interactions with, you know, for identification and authentication. In this case the technology we use is very similar to what you might experience in everyday life when you use single sign-on services. If you've ever gone to a website and you see a little login box that says something like Login with Google, Login with Facebook, Login with Twitter, it's the same underlying technology. Right?

Now one thing I want to make sure you're not walking away from, we are not talking about access to this using your Facebook credentials, all right? No, that's not it. I mean same technology. We need a little bit more than that here though. All right? But it does, you know, lead to not just an assumption but a definition of a couple of prerequisites.

In order for this system to work, this system or this network of service providers, authentication providers and, you know, authorization services has to exist and so there has to be some, some people call this accreditation process that happens, some type of a vetting process, you know, some type of a setting up of entities to provide these services.

The identity provider functions in particular which, again, in the (OOFF) and Open ID connect sense is the combination of that authentication provider and authorization service, as you'll see on the screen, is going to be responsible for managing credentials for requestors or end users. And in this context this is kind of similar to what, you know, Google and Facebook and whatnot do.

The end users have a relationship with these providers somehow and in the context of an RDAP service this could be a community of interest like yourself for example. If you happen to have a concept of membership and you understand, you know, who your members are and are willing, you know, to give them a credential and then vouch for their identity for when they ask for access to a resource, that's the concept that we're talking about here. Right?

Now there is software behind this, right? And the typical interaction begins when a human, an end user, uses the client application to perform an RDAP query, and they do that by sending their query to this ICANN RDAP access service. And they also send, you know, some indication that they want to authenticate them self.

Now the RDAP service doesn't know who the entity is and they have no information, a priori, to do any form of identification authentication but the identity provider does, right? And because there's a configuration step that took place ahead of time so that the relationships between these entities are well understood, the RDAP access service basically vectors the end user over to this authentication provider who pops up a little web interface. Again, very similar to what you see if you say sign in with Google. The next thing you see is a Google login page.

You fill out your user name and password, you know, push a button and then the next thing you're asked is I need your consent to share a certain amount of information with, you know, the resource that you're trying to log in to. And this is where this does get to be important because in this context there's going to be a need to develop certain attributes that are associated with these identities. An these attributes have to be the kinds of things that can be used to make an authorization decision.

So things like what - who are you, what is the purpose of your query, what role are you performing, in what legal jurisdiction do you reside? And that's by no means an exhaustive list and it's not something we tackled directly. There's a lot of policy work to happen in that space in particular. We just note that the model accommodates it. Someone else has to figure out the hard parts. But anyway.

So something pops up and the human is asked, you know, tell me these bits about yourself and give us your consent to share that information. Okay?

Push another button. Information is sent to the RDAP service again and now the RDAP access service knows that the identity provider has validated the identity of the end user but we still don't know if they're authorized to see what they're asking for. All right? And that's where another query gets sent to this authorization service.

And right now the model we are envisioning is that ICANN would operate that service. However, the model is flexible and if a third party is an appropriate policy decision to operate such a service, it accommodates that possibility. Right?

And the way this works is it sends the same sort of information, the query, what are they looking for and all of the associated identification information, including all of these identity attributes, and there's some magic sauce on this particular thing that maps this information to policies or access profiles, authorization levels, you know, use whatever term you want to use to describe it, but think of it as a mapping.

And then ultimately an answer is returned to the RDAP access service. Is this person authorized to see what they're asking for or not? And in the case of not, they basically get a, "Sorry, no thank you. You know, you can't see this type of a response" back. But in the case that they are authorized, the RDAP access service then submits a series of queries to the registry/registrar RDAP service as appropriate in order to form a complete response.

All right, remember, not all registries are thick, right? The - my employer for example still currently operates a thin registry. That's VeriSign by the way. And so in order to collect a full response for a .com domain name, you have to query both VeriSign and the registrar of record for the domain in question for example. All right?

The RDAP access service gets all of this information, forms a complete response and then delivers that to the client for display to the end user. That's the model in a nutshell, big nutshell, but.

Ram Mohan: Thank you. Let's move to the next slide. We are in - currently in the community input part of the - of our work. We're looking to get your comments and your suggestions and to also identify, you know, what we've got wrong and perhaps if we've got anything right what those might be.

And our intention is to spend the next few weeks discussing and looking at the feedback from all of you in the community. Our expectation is in the middle of April that we will meet face to face one last time to finalize the technical model and we will ship that model, we will publish that model on the 23rd of April, next month. And at that point the technical study group will conclude its work.

We will be done and the model will, you know, hopefully live on or maybe it won't but either way what we know is that it will morph. It will take on other shapes and things like that. But hopefully what we have left the community with is clarity on, A, how to go about doing such a task, and, B, clarity on what kind of a framework might actually be an appropriate applicable and potentially feasible framework for doing this access to nonpublic technical registration data.

So with that I'll open it up for questions. I'm not managing the queue so would you like to manage the queue?

Stephanie Perrin: I'll be happy to manage the queue, if you like.

Ram Mohan: I'm happy - thank you for helping me on that.

Stephanie Perrin: Yes. Okay. So first I think we had Kathy Kleiman in the back there with the microphone.

Kathy Kleiman: Thank you so much and sorry for making you turn around. And, you know, thank you of course for your work. I know many of you have been in the field and working on this for many years.

Okay so question for you. It's my understanding, and I yield to Stephanie who's much more an expert on GDPR, that this is not a mere credential. This is not like me logging in to Google so I can my calendar or logging in to Facebook to see Ram's pictures. This is a balancing test that when somebody with a credential - even with a credential, even if I knew who they are, that we have to evaluate both the requestor and what they're requesting and the registrant's rights and protections.

So how does the system know that it's a Planned Parenthood clinic that has not published its address because it's located in Texas and would be protested and people would be blocked from getting into it? I mean - and somebody requesting it is not someone who should have access? How is a balancing test built into a credentialing system?

Ram Mohan: Great question. I don't have an answer for that. Andy or Scott, do you want to speak to the - oh, you do. Okay.

Scott Hollenbeck: I can talk in general but, Kathy, it gets to the question or the thing I said about the need to develop appropriate policy here, right? We don't have the answer to that question directly. However, you do know certain things about the requestor and so one of the tasks for the policymakers will be to identify the attributes that are necessary to make an authorization decision. Right? So we need to know something about who is the requestor and what kinds of things do we need to consider in order to evaluate, you know, their I don't want to call it rights but their ability to see, you know, certain data elements.

The hand waving part happens when you try to describe exactly what that is. I mean ultimately it becomes a bit of a mapping. In terms of automation, right,

there's going to have to be something that says here is the appropriate policy, you know, based on what this particular data element is and who is asking, what decisions have been made in terms of the authority to release that information, not so much a technical question because the model can implement whatever policies are appropriate.

So I said I can't answer the question directly. You talk like a Planned Parenthood address.

Kathy Kleiman: One answer that's simple is human intelligence. It's human intervention to evaluate dangers to the registrant or ask the registrant. But I'm not sure about automation here.

Andy Newton: Okay. So one of the things this model does is it allows for the breaking apart of authentication and authorization and distributing that to third party actors as is, you know, desired by policy. And one of the things, and the use case that you just gave us is let's say that one of the identity providers is the FBI, I'm not saying that is going to happen, I'm saying that's just a possibility. So would the FBI be allowed to have access to that data if they needed it? I don't know. Probably.

But that doesn't mean that anyone who has some sort of credentials can get into the system. It's up to whatever the policy is that's determined by authentication and authorization and if policy allows for those things to be distributed to identity - third party identity providers that are accredited in some manner, I don't know what that would be, then you have some - you have - you can have or the system can have credentials that can be trusted in a manner that we're kind of used to in society.

Ram Mohan: Thanks, Andy. Who's next, Stephanie?

Stephanie Perrin: I think Anriette was next and then Michael. Oh, Raoul. Yes.

Anriette Esterhuysen: Then you, then Michael and then me.

Raoul Plommer: Raoul for the record. So why isn't ICANN allowed to know who is making the query?

Ram Mohan: Could you bring the schematic up a couple slides earlier?

Scott Hollenbeck: Scott Hollenbeck here. Actually ICANN does know who's asking. They can't verify or validate that information though. So for example, if you think of a credential, you know, looking something like an email address, right? That's what the RDAP access service is going to see, okay? And then it's going to depend on this authentication provider to determine the validity of that credential. I mean is this person actually who they claim to be, right. Interaction and exchange of things like a password or verification of a digital certificate or whatnot. And one of the things that comes back in response or successful completion is something called an ID token

Man: So I understand that they check the validity of that requester. But isn't ICANN allowed to know who is making the request?

Scott Hollenbeck: Well that's just it. They are. The information that comes back in the identity token will include certain attributes associated with identity. And if policy, for example, determines that you must provide your full name, your date of birth, you know, the names of your three children or whatever, all that information will be available and visible, you know, to ICANN or the operator of access service.

Man: And just to clarify, our model doesn't predicate or doesn't require that the identity must or must not be known. It is dependent upon whatever the policy is. The model allows for information to be brought back. If you see those arrows, information can be brought back. What that information should be is outside of our purview.

Woman: State your name and…

Michael Karanicolas: Michael Karanicolas for the record. Thanks so much that presentation, that was really interesting. And especially from - I'm someone who's - as far as you can get from a technical background. So I appreciate if you can bear with me. With regard to the system for assessing credentials, does that allow for any centralized way of determining if credentials have been compromised? So is there any centralized methodology for that? Like, is there a two-factor requirement? Or would that depend on the individual policies or safeguards of the external authorizing bodies?

And as a second question -- if that's all right -- kind of related to what (Raoul) was mentioning. You mentioned transparency earlier on. Does your system allow for auditing of use? And how might that work? Or is that straying into a policy question?

Scott Hollenbeck: Thanks. (Andy), I think you can handle both.

Andy Newton: Yes. So if the identity provider requires for their users to have multi factor authentication or use digital certificates, then the system can accommodate that. So that's how you would you would allow the credentials to have high verifiability. I don't know what that word would be, but anyway. But that's how you would place strong confidence in the credentials.

As far as auditing and transparency, we did discuss that, we did think about it, and we do think it's important. The - so we discussed logging of these queries both at the ICANN gateway and that the contracted parties and the ability for reconciliation. And we even have recommendations about publishing usage reports so that the broader community can see what's going on with the system.

Man: So, bad things happen. You asked about compromise as well. So (unintelligible) compromise. So I think that the key to that is having audit logs

available. Not of the data itself, because of course that would - the system doesn't provide for the data to be stored at that central access service (unintelligible) that's just a passthrough A proxy, if you like. But there will be logs -- use logs -- available. Or, it is possible for there to be logs.

So I think yet I think if you - if credentials have been compromised, you'd have to look at the logs to identify what the compromise ones. So that would be evident from the usage logs and from a reconciliation process. Like with any compromised service. But of course the hope is to build strong authentication in the first place to reduce the likelihood of that. Does that answer your question okay?

Woman: (Unintelligible).

Juan Manuel Rojas: Okay. Thank you again for your presentation. This is Juan Manuel Rojas for the record. Just - and forgive me if you already answered this, but just for clarify. First is like requesting ones. So first is - okay, you are talking about non-public data, right? But what data are you classifying as non-public? Right? This is the first.

Man: Let me, let me…

Juan Manuel Rojas: Wait, wait, this is…

Man: May I respond to them one at a time?

Juan Manuel Rojas: They short, they're short. I think that they are, it will be short. So (unintelligible). Because they're so (unintelligible). Another thing is, how will the data be managed? And where will be store finally? And what about how do you manage or how would you thinking in this model about the data quality management in this model? How you will be do that? That's it. Thank you very much.

Man: Thank you, great questions. The only reason I was suggesting that I respond to them one at a time was because my memory is very bad. What was that?

Juan Manuel Rojas: The first one, it was what information are you classifying as non-public data?

Man: We're not. What we're doing is we're - we recognize that there is going to be public data and there's going to be non-public data. The system allows for access to non -public data. What is non-public data is something that others is in this ecosystem are going to define - policy makers and others are going to define. Once somebody defines that there is nonpublic data, here is a model for how to access it.

What was the second question?

Juan Manuel Rojas: It was about where it will be this data stored finally?

Man: Great question. In our model we have that ICANN (unintelligible) backs the servers is talking about is designed to be just a passthrough. No data is actually stored there. And the way we envisioned it and the way we are suggesting it, data stays - or the data is authoritatively stored. So in other words here -- in this model here -- if the data is at the registry or the data is at the registrar, it just stays there.

Data quality management. Well, I mean, the responsibility of management of that data, you know, belongs to various parties who are under contract and who have requirements. This model doesn't actually speak to - or seek to improve data quality or anything like that, or even measure data quality.

What it does speak to is that the kind of access that occurs and that the - from the various components - if there are service providers for those components, what once we are recommending is that there be service level agreements that exist, so that - including, you know, for ICANN if it stands up.

And (unintelligible) backs a service We're recommending that there be service level agreements, that there be visibility into updates or accessibility of those services.

And in addition to that, we're also recommending that ICANN -- should it set up such a service -- provide some kind of a transparency report -- even in cases of requests that have come through that may be confidential or must be required to be kept secret, et cetera -- to still provide a transparency report that reports on the metadata associated with those queries. So those are the recommendations that we're making.

(Andrea): And that's me and then it's Louise I think next. And (unintelligible). This is very cool, thanks very much. This is - I wasn't expecting to actually learn about this. But it is very cool. A few questions. I think it's very clever to separate the identification and authorization processes. So that's, like, a - I think it's going to be a really useful principle. And I think the back-end policy mapping is probably going to be much more complicated than the technology.

And I think (Cathy)'s point about building in some form of human verification should be taken seriously. And - but then just a few questions. How do you prevent rigging? How do you prevent repeat request? That then actually become - because if you've worked out how to get authentication, that could be then abused. So how do you how do you prevent that?

I think there's (unintelligible) in human rights there's these concepts of necessity, proportionality and legitimacy. And I think those could be quite useful to look at. There's even a set of principles called necessary and proportionate, which were developed for surveillance, but they might actually - I haven't looked at them for a while. They might have some useful principles for the ICANN community.

And then and transference reports, that's excellent. But what about disclosure to the entity about him the request has been made? When will that happen?

At what point? With government request to intercept or monitor communications, in the US there's a time frame that you're informed so many months after your communications has been - you know, if you've been bugged. You have to be informed a certain period after that.

Other countries you don't necessarily have that. But that's generally considered very important, that if a request has been made about any of your data -- even if it's a legitimate law enforcement request -- that you are informed about that request So does that information disclosure happen in the transparency report? Or maybe on a one on one basis, et cetera.

And then remedy. How will this - I mean, obviously, that's not your role. I get that's policy. But will there be mechanisms for people to file complaints if they feel that requests were not legitimate?

Man: Thanks, (Andrea). Some great questions. On remedy, we have in our report we have some considerations on remedy. What we've said is that we recognize that that is an important thing and work has to be done. We didn't believe that that was our work necessarily, but we pointed out - in our considerations section we have pointed that out. Similarly, the question that you asked prior to that, that's a policy question as well. And again we've taken note of that.

You know as we were deliberating there were probably at least, what? Half a dozen, maybe a dozen such issues that popped up. And what we've done is try to have high fidelity in simply recording them and placing them in in front of the community. So, you know, transparency, auditing, logging, tracking - all of those, we have some of those in the system requirements, but we also have made note that what should be done with it or how should be done are things that it should be discussed and responded to in the community.

However as a technical model, what we're saying is that the system should treat logging and auditing as a first-class requirement. It should not be

something that is an afterthought. Did any - (Andy), did you want to answer some…

Andy Newton: Yes, I want to address the notification part. I don't think we actually took that into consideration. At least my faulty memory didn't tell me we did but we should probably add that. I will note that notification has a danger if you want ICANN to do it because that means -- especially if it's doing it well after the fact -- ICANN's storing data that we would otherwise probably not want them to store. The - it may be better from a policy perspective to have the contracted parties do that. But, yes. We - I do think we need to add that to our report.

Man: Thank you. This is Scott and then Benedict.

Scott Hollenbeck: Sure. The first part of your question about rigging. Let me handle the technical piece. That's a little bit easier. RDAP is a web service, right? And so all of the tools and techniques that are available to server operators that you can use for things like rate limiting and whatnot are - they're at your disposal. And one of the other neat things about this technology is that these tokens that I describe, the proof of authentication is time bound, right?

And the validity period of a token - you know, it's a policy matter. So the identity provider when they issue these tokens will tell you when it expires. And it may be that these are very short lived credentials, such that every - it could be, you know, every time you do a query you have to reauthenticate yourself And I know there are people who are not going to like that. Others want very, you know, long lived credentials So the right answer is probably somewhere in the middle.

But it's those kinds of things. Oh, and of course, is anything you can do on the server itself. The RDAP access service can actually implement code this says - you know not just rate limiting from an HTTP perspective, but hey,

you've asked me X number of questions in the last 30 seconds. Slow down. So any number of techniques.

(Andrea): I think Louise was next. And can I get…

Man: Sorry, Benedict also had a response.

(Andrea): Oh, sure.

Benedict Addis: I was just going to say this is really cool. Because at the minute, you know, we've got a proposed model, we suddenly start to think, well what about the rate limiting and what about, you know, where can you break this? You're thinking like hackers, which is brilliant. I think that's probably quite a line. I actually had a question for you but perhaps you'll have to mull over, which is - and we've been - because we've been hearing this from registrars as well.

At the moment, ICANN, that middle part is the gateway, obviously. It's not storing data, but it is the sort of arbiter. The registrars had talked about wanting to do manual approvals as well, so this is a concern that you share with them. And I'd be interested to know (unintelligible) you just have us consider when you're recommending policy where that manual - if you see that bit needing to be a manual step, where that should live.

So in this - again, this is only a proposed model, but in this model it can accommodate that as we've heard. But where should it live, and have a think about should it be applied to identities? Or should it be applied to requests? And obviously there's a trade off if it is requests, because then you're - there's leakage of data. So all of these things come with - they're all technically feasible but think about where you want them to sit And perhaps how time - whether that approval is time bound as well.

Man: Thank you

(Kate): I'll just jump in here. I think we should let people know that we have nine members of this group on the EPDP. Maybe people could wave their hands. And those are the ones who are fairly aware of the answers to some of these questions, like where's the data going to live? At least from our perspective.

So and if I may, also, grab the microphone for a second for the benefit of our members who haven't heard me talk about this before, University of Toronto has a - has been funded by the Office of the Privacy Commissioner to investigate an access model I'm looking at it from the perspective of policy and procedures, and what is necessary under data protection law. So we're kind of the mirror image of what's been going on at the - in other words, we're doing all the hard questions. And you guys got the easy bit, right?

And we did have a workshop in Barcelona, which is still up in the website. Verisign came and talked, and so did Microsoft and several others. And I'd be happy to talk about it more later. But in full transparency, I've also been talking about it. I know that the board has been talking to data commissioners. I don't know who. I've been talking to the Berlin group, because they are the technical study group that looks at telecom issues, and reporting on you know the kind of project we've got going. Thanks.

Man: Okay. Thank, (unintelligible) (Kate). So there - back on that assumption slide, I think they assumption - it said you given the assumption that you would have one unified ICANN access (unintelligible), which is a pretty big assumption I mean that's a big change that has not - you know, the APDP has not discussed. So that's clearly I- a very large assumption. And if it was not, is there any way in which -- for example - a registry or registrar in this model has any say about whether or not to accept a particular request? Or has that basically all been sort of outsourced to ICANN.

And also in that assumptions -- a sort of related thing -- the assumptions there is that they - did they- the people we - the initial bridge use liability, I think, was that the term? Your liability? And so this has been a subject of

some discussion. Because IANN, of course, is not going to act - does - it seems to be fairly adamant it's not going to accept any sort of have legal responsibility in the sense.

So they must be talking about in some way procedurally it will reduce the likelihood of liability or something like that. So that assumption may not be true. And certainly the issues about who is acting as data controller and stuff are very unclear in the EPDP. So this assumption that of course ICANN will provide a centralized service, and everyone will want to use it, you know, jumping ahead a fair way in the process.

So if that was not the case, well - two questions, really. If it is the case does - it that ICANN is - it's all going for a centralized ICANN service, do registries, registrars in this scenario have any sort of say in this? Could they sort of go, well, I don' care what ICANN says, I'm not accepting, you know, I'm not accepting, say, law enforcement request from Saudi Arabia. And the - you know, for a registry in some completely different country.

And, yes. And is - would they be able to do that run their own sort of RDAP on a specific thing? And I certainly know there are some services that might - may or may not be providable if there's a central service, in terms of - I mean, I know a law enforcement and security are very keen on the idea that they can do this - you know, you can search for any matching credentials across multiple providers.

And I and I totally understand why, you know, people want that. I guess that's probably enough things lumped in together and just sort of a question yet as anyone, yes. If anyone care to comment on that basic can you desegregate or aggregate sort of (unintelligible)?

Man: Thank you could you take us back to the assumptions please? Thank you. So ICANN org Euron has asserted that there is a belief that ICANN's desire with the unified axis model so to speak is to reduce the GDPR liability of gTLD

registrars and registries. So we took that as an axiom, right? So we took that as an axiom, and we stated it as an assumption.

Now, we did the work and as we did the work what became apparent to us was that - and you'll find this again in our other considerations section of the of the report - be saying that in that section that we actually can't speak to whether this assumption is true And we also - and therefore we exhort the - those parties -- the registries and registrars -- to exercise their own discretion and to engage with whatever counsel they need to arrive at their own conclusion, right?

So that that's a spectrum that that we traveled. First making a note of the assertion and the observation. And then later on saying we can't independently verify it, but certainly the parties that have liability ought to go and verify it themselves. Scott, did you want to respond to the rest of the questions?

Scott Hollenbeck: Sure. In terms of the first part of your question you know what is the - I mean, can the contracted parties say no? Well, the model gives them the ability to do so, in that while we get queries from ICANN, we also will receive the information that identifies the original requester.

Louise Marie Hurel: Right.

Scott Hollenbeck: Right? And so in theory, you could say that if the query violates some local law or policy or whatnot the contracted party could very well say no I declined to do so. Okay? All right.

Andy Newton: So I also want to talk about that first of assumption where ICANN is the point where all the requests, or all the queries go. That was kind of the -- as Ron said -- kind of something that was handed to us. It does turn out there are some technical benefits to it though. And one of those technical benefits is that with ICANN acting in the role that it is, you don't have to invent some sort

of policy language and a policy distribution mechanism to send that to all the contracted parties.

So there is - the system itself becomes much less complex in regard to how you distribute policy and how you even talk about policy amongst the actors in the systems. The other thing is, it also lowers the threshold for what the contracted parties have to do as far as knowing who to trust. And the model we put forward, we said the contracting parties only have to do mutual TLS with the ICANN RDAP gateway. And that makes it much simple for them. So there are technical benefits to doing them.

Man: Thanks. Benedict, did you want to add something? And could you please take us to the proposed a schematic diagram?

Benedict Addis: Okay. So another advantage that this group might be interested in is that access service being centralized -- and again, we highlight, just that access service -- is that that you get centralized logging, which facilitates the production of a transparency report. So that's - it's - and that can be reconciled against logs seen out at the registries and registrars -- again, can be -- to make sure that everybody's playing straight. There's no discrepancy between the center and the nodes.

And you can accomplish that without storing or retaining any data at the center. So it seems to me to (unintelligible) quite useful. But it has some (unintelligible) has some benefits.

Woman: Can I get in the queue now, do you think? I've been doing a terrible job in the queue. No, (unintelligible) is the queue minder, okay? I'm in the queue next.

Louise Marie Hurel: Perfect. Hi, Louise, for the record. I don't know - I've been - well, I'm trying to follow all through the technicalities. But I think (unintelligible) covered most of my concerns with kind of like the misuse and already kind of like trying to break the system. But on the other hand I was wondering about the metadata

that is produced by each of these authentication providers at the authorization service. And all of them will have logs.

So in terms of data storage, even though I completely understand that the data is being stored by the contracted the parties, they are the ones providing the central data. But in terms of auditing, I see that there must be some kind of metadata coming from these servers that will allow it to be audited. So what about the time which this logs will be available? And what is required or what is necessary - I know this is more of a policy discussion, but I was just wondering, because you said, no there's no data.

But, like, I'm thinking about the metadata and the trails that kind of like authenticates the identification and the request. But at the same time how it serves different purposes, you know? I'm not sure if I made myself clear, but yes.

Man: That is an excellent point and I'm not sure we've actually spent much time discussing the specific metadata pieces and storage and logging and tracking of that at the authorization service or the authentication service. Scott, did you want to respond to that?

Scott Hollenbeck: Just a little. And only confirming what (unintelligible) just said. We talked about it a little bit. Remember, these are all Web services. So for one, Web servers log HTTP requests, right? You - that kind of happens by default. And typically with the way RDAP is structured you're going to see certain elements in the path segment such as the domain being queried. And with the way it's currently structured you may well see at least the format of the credential being used.

Like, this thing that looks like an email address. And sure enough, if you collect enough of this information you've got a very interesting data collection, right? We have simply noted in the report that yes, this information is logged. And there's going to need to be policies developed around what happens

there. Now, the layer above that, though, again, we talked about it and we didn't really come up with any solid recommendations.

So for example, above the HTTP layer, these services are going to be performing certain functions. One could assume that they will be additional logs created, right? To verify, you know, the thumbs up thumbs down aspect of, you know, how these things are happening. No, we didn't touch on that in any great detail, other than knowing that it's possible and maybe we'll come back to it later.

Louise Marie Hurel: Okay.

Man: There's an (unintelligible) that talks a little bit about the - it's about privacy, but it talks about different levels of, like, yes like being anonymous versus being unobservable. Like, if you purge everything, it becomes unobservable. So that's something we would probably want to avoid keeping track. So there will be some amount of data that will be collected.

Louise Marie Hurel: Sorry, just a quick follow-up. So in this - what I'm hearing perhaps is this should be delegated to kind of like the policy discussion? Because it's very much entrenched as to how the technical system is being built. And even though the decision on how long should a person be able to access the logs is a very, like, policy-oriented discussion, on the other hand I do see that the embeddedness of the technicality of what kind of trail is being set by the system seems like a very technical in nature. So, yes.

Man: So, I agree with you on that. I think there are both components of it, and, you know, when we sit down and we meet tomorrow, I think we'll end up discussing that piece -- the technical piece -- and reflecting on whether we should add something to the requirements section as to what should happen at the authentication and authorization services.

Benedict Addis: And remember what Scott was describing was logging of the requestor, so who's asking? And I'd be really - and encourage you to think really strongly about splitting those apart. Who's asking and who's being asked about? And what your desired policies about both of those.

Man: Right. So I think those are really important things and those belong in the policy part of the discussion. And what we'll do is, you know, reflect upon what you're saying here on the technical side to at least arrive at some either conclusions or to provide some thoughts on whether the model can accommodate, you know, either metadata storage or metadata purge, or all those kinds of things. Thank you.

Stephanie Perrin: Just a - it's Stephanie Perrin here. Just a couple of administrative notes. Theoretically we stop at 6:30. I'm sure staff would like us to stop at 6:30. GNSO counselors and APDP members have a very short appointment way the heck back at the main venue. So we're going to start losing all of us pretty soon. And it's clear that we could go on for hours talking to you, because there's an awful lot of policy issues that we'd like to see sorted.

You know, your credentials have got to be unlinkable, you know? You have a lot of data retention that you have to do to meet data protection requirements. Jurisdictionally, one of the reasons that I have been saying we need this data trust to be not at ICANN - there's political reasons, there's jurisdictional reasons. We can't have it in the United States, or we'll just be up against an adequacy decision very shortly.

And even though I understand the data is being held at the registrar registry level, we're pushing for thin registries uber alles. By the way, kind of a policy position we have. Even though that's the case, the deterrent - the decision making is still happening. The processing -- which is an administrative decision -- is happening at ICAN. So you're still stock. And I know you've probably all heard me say this 46 million times.

But ICANN is the controller and hasn't accepted that mantle yet. And normally in a data protection analysis, the first thing you do is sort of who's controlling things. And then you - you know, then you map it from there. We're going backwards, and it's extraordinarily frustrating. You know? Yes. And in terms of the liability issue - and I'm not a lawyer, I should say. I'm going to do the caveat here. I'm not a lawyer. I'm only a policy person.

But there's two aspects of this - and once again we're using in terms that not everybody's understanding here. Obviously if you build any kind of compliance with data protection law, you are reducing liability. And by building in good security, you're reducing your risk, right? So you're doing the appropriate things. Your liability is lowering. We can all agree with that. But the controllers, you know, the registrars and registries don't get off the hook.

They still (unintelligible) controllers under this scenario, or joint controllers. And they need quite a complex mash of data processing agreements to sort out who's liable for what? Who's retaining what, you know? And that's - it's going to be better than now, obviously. But it's a non-trivial set of tasks to sort out. And that's not a question. It's a rant.

Man: Thank you, Stephanie. I think -- and I speak for the (unintelligible) -- that we largely agree with all of the things that you're saying. This is a multilayered and very complex problem. You know, earlier in a session that we had, we had one of our members asked quite specific questions on how long will it take -- assuming policy is all done -- how long will it take to implement this?

And really the answer is the technology is not the bottleneck here. The feasibility of building a system - you know, you've seen this demonstrated. Here is a proposed model. You know, we started our work formally I think in November - late November, something like that. And here we are in March with something that is at least ready to be digested, right? So technology is not going to be the problem space here.

The final thing is we -- at least the way -- I agreed with (unintelligible) when he asked me to set this up. I said I'll only do this as a specific time to project. Because the subject matter here is a very large subject matter. And the work that this volunteer team has put it has just been stellar. But it's also been all consuming for a little while. As all of you were doing with the EPDP and other things, right? So we wrap up our work next month.

But I certainly expect that there will be TSG2 and TSG3 and, you know, all the way to TSGN of some sort, right? It won't be me, certainly, in the coordination role of all of those. But I certainly expect more work is going to be done And I expect that this model will evolve as policy answers come through. Clearly these answers - the technical model and the framework underlying that will have to evolve.

But we're really hoping from the TSG that both the way to engage, to arrive at technical design and technical solutions, that we've been able to show a way to do it. And that the framework that we've come up with might be a framework at least some of whose components might last. For instance, we've said in in this model we have has intentionally and explicitly broken out - not just this identity provider, but also that there be an authentication provider authorization service broken out as components.

There's no requirement that they have to be in separate places. But we think that from a model point of view it's useful to consider these as components that in some cases could be distributed or delegated et cetera.

So with all of that, thank you for having us here. Thank you for your questions. It's been delightful. And look forward to your - if you have further questions or if you have more things, please write to us. We're going to consider all of those things when we do our deliberations.

Stephanie Perrin: Well, and we certainly have other members of the APDP who weren't able to be at this session. I'm sure we'll have further comments and questions. I'd

also like to just clear the air because Benedict was giggling at me. I had a minor tantrum. That's only a minor one. The other day on that cross-community session, it was certainly not - nothing to do with your RDAP proof. That was fine.

It was that the business community constituency is trying to represent that we have agreed as a multi stakeholder group in that slide. And I'm still cranky, as you can tell. I think it wasn't accurate at all. But it certainly wasn't directed at you folks. You did a great job. And I agree to get this all wrapped up and mapped out, it was really useful. So thank you.

Man: Thank you very much.

END