KOBE – Community Engagement Session: TSG on Access to Non-Public Registration Data
Monday, March 11, 2019 – 13:30 to 15:00 JST
ICANN64 | Kobe, Japan

RAM MOHAN: Good day. We're going to get this session on the Technical Study Group on Access to Nonpublic Data started in just a couple of minutes. Thank you.

Good afternoon. My name is Ram Mohan, and I'm the coordinator of the Technical Study Group on Access to Nonpublic Registration Data, or TSG-RD as we call it. We have about 90 minutes scheduled for this session.

In the session, we expect to take perhaps 45 of those 90 minutes in actually walking through the process that we've gone through and presenting to you a draft technical model for your feedback and for your comments and your input.

And our expectation is that this will be an interactive session. This is not the only session. We met with the EPDP folks yesterday, and later today and tomorrow, we intend to meet with several other groups as well to present what we have done.

So with that, let me just walk you through our agenda. We plan to cover these topics. Expect that we'll have plenty of time to address any questions that you have. And I believe we have folks here who are also taking care of looking at comments that are coming in from the remote participation, so we'll be able to handle that as well.

So with that, let me just begin with a little bit of an introduction about the Technical Study Group itself. To really make those, to set the scene, to set the stage for how we got started, what we're about, etc., let me turn to Göran.

GÖRAN MARBY: Thank you, Ram. Before I start, I would like to thank the members of the group for the hard work you put in for the last three months on top of everything else you do. Seeing it from the outside, I'm always happy and blessed to have volunteers doing this work.

So, why this group? First of all, we talk about something that is legally interesting with a potential technical solution. So there's a couple of things I have to repeat, and you've probably heard me say it about 200 times before.

The law of GDPR is very specific about the role of the data controller, data processor. The ones who take in the data, have the data and make decisions about the data is the one who is responsible under GDPR.

In our world, that becomes the contracted parties. ICANN Org as a legal entity doesn't have the database. I know that this [has] come as a surprise to you.

So with that specific, it was fairly obvious for us a long time ago that it's very hard to do any sort of unified access model with one vehicle, because the contracted parties as individuals have that legal responsibility.

What that means is that even if we had a policy about it – I see Mr. Steve Crocker arriving. Can I point out that Steve Crocker's been working for me for the last three months? Hi, Steve!

STEVE CROCKER: Hi!

GÖRAN MARBY: It just felt so good to say that. I'm sorry. So that is sort of the assumption to this one, is that the contracted parties have to make individual decisions.

So we started looking at different solutions, how to diminish the contracted parties' legal responsibilities when it comes to WHOIS. And I don't know if you remember in Barcelona, I received a letter from I think almost all contracted parties, who said, "Why don't you go out and see if there are any potential solution to diminish their legal responsibilities in order to create the unified access model?"

Because if we don't change the current interpretation of the law by the DPAs or come up with [something that changes the] game plan, as you know, it's really hard to do a unified access model, because we can't enforce from ICANN Org something that goes beyond or above the law. The law always drives ICANN.

So the background where we started this discussion was really that we had conversations also within the European Commission about different alternatives where we started looking into if ICANN Org legally could become the place where you go and ask the question. And the idea was set up, which is sort of based on RDAP, that you come to ICANN org with a question, the question is then transferred in a secure way according to GDPR principles, to the contracted parties. And the only one who can answer that question is the contracted party back to ICANN.

That seems very easy on the surface, but we also realized that we needed real technical knowledge to work on that model. And instead of us inside ICANN Org just sitting down and doing a model, I decided in my wisdom – which I think was wise – to ask Ram to collect a group of people with high technical skills to look into that.

And now of course the question comes, what happens now after this presentation? First of all, I'm looking forward to your input o the solution that they have worked on.

Now, after that, what we're going to do, because this is a part of a major system, this is the sort of exchange point. On the side of this, there has to be someone who recognizes who can ask the questions. Well, [inaudible] general terms called accreditation houses.

And there are several accreditation houses out there. I can mention for instance Europol. I know there are also parts of this community who are looking into different ways of building accreditation houses. We suggested a long time ago WIPO as a potential one.

And the idea is to connect those accreditation houses who validate who are the requester, who validates the questions through this [mechanism, asks] the question and that goes to the contracted parties.

The way we've done that work is sort of built that together. Our intention is go back to the data protection authorities to ask the simple question: does this diminish the contracted parties' legal responsibility?

If the answer is yes – and here we're using the word "guidance," and if any DPA would listen in to this, the guidance in the European perspective is actually legally binding in that sense, that when a DPA gives a written answer to a question, that holds up as a part of their decision-making process. So it's stronger. That's the word. that's what I think. Data protection authority cannot say something without in that legal context.

So the idea out of that is that it actually gives the ICANN community, if this happens, the opportunity to create policies for unified access model. Who should access it, what do we think about the privacy and access to information?

The intention of this is not to take over the policy work within ICANN, it's actually to be able to inform the ICANN community about the legal possibilities for unified access model, or whatever we want to call it.

I want to caution you that if you go into the data protection authorities' webpage, they will say, "We're not a consultancy, you have to go do your own work." To receive legal guidance from any data protection authority, it's something we have to work very hard with, especially together with the European Commission.

We were one of the few who actually received any guidance whatsoever during the last – what we call the [calzone] process.

And for the record, I'm not allowed to name any project more inside ICANN.

So, that's one of the sort of problematic [inaudible]. When we went into the temp spec, we actually had legal guidance from the DPAs. And remember what that legal guidance said. It said that [we had the] right, we could collect data. Didn't specific [which,] but it said we can collect data.

It also said that they accepted we had sort of a covered model with some information public and some information nonpublic. That was very strong and very important guidance for us. Without that, I think that the DPA [inaudible] work has been much harder.

Now we're in a situation [where] we don't have that legal advice for phase two, so that's what we're trying to address. I open the floor for any questions.

No one? I feel very lonely suddenly.

UNIDENTIFIED FEMALE: One question. Is reduced liability enough? A lot of companies are going to invest billions. Is it enough to say that they there's just a possibility that their liability will be reduced? Can't we get a more precise answer of what the DPIAs categorically do not accept?

GÖRAN MARBY: Yeah, I wish. That's up to the DPAs. The law is – the GDPR is interesting. I've been a regulator, and actually, the GDPR is technically interesting. I think I have a bad sense of humor, so I once called it a [motherly] law, because when I was a teenager, my mother used to say I can go out if I behave. And then I went out and behaved, and apparently, she had another view on behavior than I had.

And when someone told her that I misbehaved, that's where I came into trouble. And that's actually the law. The law says that you should do what you think is right. As long as you can explain it, it's okay. But if it's not, we're going to come after you.

So it's open to fairly big interpretation by yourself. And that creates one of the problems of the law, is because the individual contracted parties actually have to make that judgment by themselves. And ICANN as a legal entity through our contracts has a big problem enforcing that to the contracted parties.

RUBENS KUHL: Rubens Kuhl, nic.br. I have a comment and a suggestion. The comment is that while one of the motivations is reducing liability, there is a potential for this model to reduce operational costs in addressing the unified access model. So there are incentives both negative and positive that could play here.

But my suggestion is that ICANN doesn't try to make this model mandatory. Anything that is compelled usually gets pushback, and "Why is this? Why is that this way, that way?" But if something is optional but has strong coverage inside the gTLD space, it might be stronger than something that's compelled, exactly because these are not by force. That's an idea.

GÖRAN MARBY: [inaudible] the second thing really is that we already in ICANN in our contracts have what we call waivers when local laws supersede our contractual obligations. And you might say that this is a very big waiver. In fact, it actually started out in 28 member states plus the EEA countries and so forth.

ICANN is not and should not be in a position – we cannot enforce if the local law says anything against it. And I'm saying this – I heard yesterday that this is just a dream. Yeah, but sometimes we have said that we are going to try to do this.

Do I give it 100% probability? No, I don't. But I thought it was important enough to investigate the possibility to diminish contracted parties' legal responsibility. And of course, if they accept that – because it's going to be something they accept, because that's also according to the law – the DPA's advice has to be legally that strong, that they can feel confident that someone makes this decision for them. So it's hard to [inaudible].

But on the other hand, ICANN as an institution is a fairly voluntary arrangement. The contracts we have with our contracted parties is based on policies made by the community through a bottom-up policy process. So in a way, we are in a consensus model, and consensus means that you accept. But I get your point. But we're not a government.

UNIDENTIFIED FEMALE: I have a few questions. I'm sorry, I was not part of the EPDP so I don't know exactly if these issues were dealt with. But I would like to know two things, actually. GDPR gives end users more access to more responsibility over their data. Can you tell me if for example an end user, an owner of a website who wants people to see his information for example, in that new model, would there be box saying information is public or information has to be kept private? That's the first question.

Second of all, I heard someone from a registrar say yesterday to a lady who was requesting access for consumers who wanted to check the authenticity of a website that unfortunately, since registrars could not make the difference between an individual and a company, that they were not going to make that distinction at all. Is it true, first of all, that this distinction between individuals and companies will not be made in the new system?

GÖRAN MARBY: You are asking questions about policy, which is in the PDP, many of those questions. So it's a policy question, and neither me or anyone else at this table are involved in the policies.

I draw a very strict line with that, because I think that those questions should belong [in eventual policy work and policy, in the] one or two. The technical solution has one specific target, and that's see if we can diminish the contracted parties' legal responsibility.

I accept the fact – which makes your question very smart – that it could be so that the legislators, the DPAs, could have opinions about the existing PDP, which could have an effect. But we don't have that answer yet.

The legislation is also fairly new and there are very few court cases to [inaudible] what is private data and privacy data. When it comes to your first question about if it's possible to do a total opt-in if you want to, I actually don't know the answer to that question.

BENEDICT ADDIS: Hi. My name is Benedict Addis and I was on the EPDP as well. So, speaking purely from the perspective of the EPDP and not as a member of this technical group, I can tell you that we discussed the idea of opt-in by the registered name holder to have their details published, a lot. This was under much consideration.

The answer is probably yes, that will be possible. But at the moment, it is not. So it's in people's minds very much. An answer to your second question about the legal-natural distinction, there are a number of reasons that's difficult. It seems it's one of those things that at first sight seem to be a relatively easy thing to do, and it has a great deal of complexity when picked apart. For example, there are organizations that are entitled to privacy in their countries under local law. You can imagine an abortion clinic might be entitled to that.

So there is a fractal level of complexity about that question that means that we haven't said yes or no to this question in the EPDP, but we have deferred that decision to phase two. But thank you for some really good questions.

RAM MOHAN: Thank you, Benedict, thank you, Göran, and thanks for the questions from the floor. And as further questions about policy matters and other issues come up, certainly, you could bring them to us and we will act as a passthrough to send it on to the various groups that may have the relevant approaches to those questions.

So, let me just – and Göran.

GÖRAN MARBY: I'm going to be down there and listen.

RAM MOHAN: You got it. Thank you.

GÖRAN MARBY: Thank you.

RAM MOHAN: Okay. Let's talk a little bit about what we've done and how we got started. Göran gave you that background of how things got started. The purpose of the TSG, the Technical Study Group is to explore technical solutions for authenticating, authorizing and providing access to nonpublic registration data for third parties with legitimate interests built on RDAP. So that's the purpose.

Now, there's a charter for the TSG. All of this is published in the charter. If you go to the URL that you see on the screen in front of you, you will be able to access all of that information.

Now, we were very clear, as Göran mentioned and as we have tried to add here as much as possible, that the TSG will not be making any decisions or make any recommendations on policy questions. For example, questions on who gets access, even what is access, should it be called access, what data fields under what conditions should access should be given, what is legitimate interest. There's a panoply of issues that exist, and we're really glad that they exist, but it's not in our remit for the most part. We've been focused quite clearly on the technical side of things.

Now, who are the TSG members? As we said, Göran is a sponsor for this. He asked me to stand up a group in October of last year. I spent a little bit of time making the decisions, and you see some photographs here, but most importantly, you see almost everybody from the Technical Study Group here in front of us. I think only Murray from Facebook is not here. But the rest of the Technical Study Group is here in front of you.

We've been very fortunate that the work of the volunteers on the TSG has been very ably supported by an absolutely first-class ICANN Org team, and you see several of them here. There's [Elisa, there is Diana,] there is John Crain, there's Gustavo. There's also Francisco Arias who is not here with us. And we also have had fantastic help from Yvette as well as Erika on the work that we've been doing.

So, when we got started, it was quite clear to me that the way to get to real results was to have the TSG work in a consensus-driven way and that we had to be iterative in our process, and that our focus had to be technical. So that was the primary engagement model that we went about our work.

And really, the way we put the process together on arriving at a model, arriving at a solution, was the following: we first began by defining key questions and considerations. Then we identified the main assumptions. Following that, we identified use cases as well as the user journey, and then we defined system requirements, functional, operational, management, all of those system requirements. We created a mapping to the functional requirements, built some actor models, determined implementation considerations.

And when we did all of those things that allowed us at a face-to-face meeting last month, it allowed us to arrive at a proposed solution. And that was a tremendously iterative process we had, several models that came in front of us, and we were able to go through that, look through what appeared to be good or not good, and eventually arrive at a proposed solution that we're calling the technical model.

The next part of our engagement is to – and as we were doing all of this work, arriving at a technical model, looking through the actor models, looking at the implementation considerations and the requirements, several things popped up that were clearly considerations in this entire space, but it was also clear that these considerations and these observations we were making were not things that were for us as a Technical Study Group to actually act upon.

But what we're doing is, as a matter of completeness and as a matter of really good stewardship, we are making note of these observations and these considerations as they've come up, and those would be part of the final document that we publish, but all of those are really intended for other parts of the community to go and work on rather than the Technical Study Group itself.

Once we do that, the next piece is to invite community feedback. This is one of those sessions where we are looking for feedback from you, and once we do that, later this week the Technical Study Group is actually going to meet face-to-face right here in Kobe to review the feedback that we received from here to look at whether we need to make any changes or other modifications to the model that we've arrived at.

And we will then spend the next few weeks – another three or four weeks – to further iterate the work that we're doing. And our intention is to finish our work by the middle of April and to publish what we would consider our final work product, to publish that at the end of April, hand it off to you and the community and to Göran, and we will then finally qualify to get to item 13 of our charter.

So, I was speaking earlier about the process, and if you look at what we did, the first thing that we did was to look at what the key questions and considerations ought to be.

If you go back and go to the ICANN.org/TSG page, you will find the charter there, and in the charter you will find these major categories listed and questions underneath these categories. Approximately 17 or 18 questions that belong in these various categories. That helped organize our thoughts into what should we be studying, what should we be doing about these areas.

One of the first things that became apparent to us as we started to do the work was that we had to be very clear about what assumptions we were making, because certainly, if we hadn't defined some of these assumptions, some of the core pieces of our work would have just been stillborn, would not have taken off at all.

So we began by after this process, one of the first things we did was to go list the key assumptions. We've refined those assumptions as we've iterated through the process. so we began in November, I believe, with only about seven or eight assumptions, and we're up to many more than that as we've gone through the process. That's a good sign. That's a sign that we acknowledge further issues.

Now, one of the things that I want to point out is when you see us talk about the assumptions, both in the document as well as in the slides here, what you should recognize is that we are not making an assertion ourselves that these are the right things to do. what you're actually seeing us do is simply documenting that these are assertions or these are assumptions that either have been made or that exist in the space, and that the foundation of our work is based upon the assumptions that are there being true.

Now, clearly, if some of those assumptions are either not true or need to evolve or whatever, that will likely have some knock-on effect on the model itself, and I certainly look forward to being in your space, sitting in the audience, listening to the next group that’s going to go and look at how to evolve the technical solution.

So one of the important pieces of all of this is also that the validity of the assertions, our remit is focused on the technical component. If you see the assumptions that we make here and if you see policy pieces or you feel like you have to question whether these assumptions actually will hold up, we have some statements here.

Please do bring the questions up to us, but recognize that we do not stand in a position to provide any authoritative answers on whether these assumptions are appropriate or not. What we do want to know is if these assumptions that we’re making, whether we've missed some assumptions, number one, and number two, whether the assumptions that we’re making are actually completely wrong, and therefore might undermine the validity of the technical model.

So with those as a framework, let me pass the microphone to you, Steve, and ask if you could take us through the assumptions slides.

STEVE CROCKER: Thank you, Ram. The picture here is the basic conceptual picture that queries for nonpublic gTLD data are mediated through an ICANN gateway, which takes advantage of and has access to the credential that are applied to the particular query and the authentication and the authorization processes there.

there are 12 assumptions. I've referred to six of them in the parentheses on this slide, and the next slide, I'll show you all 12. The basic model is that RDAP is the mechanism that will be used, hence Port 43 access will be deprecated, that access to gTLD nonpublic data is only via this mediated access, that queries from unauthenticated sources will be handled in accordance with the policy for that, and that ICANN oversees the credential protection and validity associated with all this.

This slide lists all 12 assumptions. The ones on the top are the same ones we just covered, and the ones on the bottom are various assumptions that effectively deal with evolution and tailoring and related issues as the model gets fleshed out. So there has to be a process for handling changes in datasets and rules, it's got to match normal RDAP use and evolve to existing RDAP practices. It has to be a pilot, policy choices have to be reflected, and implementation practicalities.

This is necessarily a very terse, compact presentation. Read the report for the rest of the details. Over to you, Ram.

RAM MOHAN: Thank you, Steve. So, having made these assumptions, and if we could please move to the next slide, we then came to define a bunch of use cases. Earlier, I had talked to you about the process that we had followed. So use cases was the next part of the work that we did. Andy, I'm wondering if you would not mind getting us started up on this piece.

ANDY NEWTON: Sure. So the use cases we went through – and this was, again, iterative so we came back and refined these as we went through, but we wanted to talk about authorized users, people who have some sort of need for access to this information. Law enforcement was I guess an actor that we kind of came back to over and over again, but there were others, like security researchers, intellectual property attorneys, people like that. But they required access to their multiple queries or they need to do even single one-time use queries.

We also said that users who receive authorization online, they need to get that authorization as immediately as possible, and then again, we also had a third use case where we said there needs to be an ability for some users to have access to data that is associated with them. And we even need to support use cases where the authenticated user may not be authorized to see the data.

And finally, we talked about users who are the subjects of the data and how they get access to it. So from there, we kind of came up with some system requirements on top of that. And again, this was iterative as well. at first, we started off by looking at different components of the system, and then we kind of broadened it out a bit. But overall, we said this has got to be based on internet standards, has to support IPv6, needs to be distributed or able to support a distributed model, and we needed to use secure protocols such as TLS and other appropriate secure protocols that may be applicable to the systems that we're specifying here.

One of the things we came up with immediately was talking about a web portal for people who need to have expedited requests or one-time requests, those types of things. So we have requirements for a browser- based web portal to be run by ICANN.

We talked about authentication and authorization determination. We split those two things apart. We wanted them to be delegated if possible to qualified agents according to ICANN policy. Then we talked about how we would actually do this, and we had this concept of an ICANN-run RDAP gateway which queries the contracted parties, their RDAP servers. And we said it's got to support multiple authenticated requesters and their identities and different policies that go along with that.

Has to be able to deal with granular access to the various data elements, it has to support passing of attributes of the requester to the contracted parties, and when you get an unauthorized request, it has to know where to redirect that. In addition, we need to be able to support automation as well.

Then there was the RDAP servers run by the contracted parties and they needed to basically respond to requests from the ICANN RDAP gateway.

Let me get in some more general system requirements .we have requirements about logging and auditing where we want these queries to be logged. We need some sort of ability for data retention, and we need to have a way to reconcile the queries from all these parties so we can do audits and deal with system abuse.

One of the other aspects of total system-wide requirements was we looked at performance and service-level agreements, and we said there had to be service-level agreements for all the subsystems because without that, you never know what part of the system is breaking down. So there have to be some sort of guarantees for what's going on.

Finally, we looked at information security requirements and basically state that there needs to be an assessment of what the requirements are for that, then there needs to be a way to undergo audits and provide the auditing information to those who request it. and finally, if there are breaches, there need to be ways of reporting those breaches.

And then lastly, we looked at organizational controls and we said this needs to be governed by a business continuity management program and we need to make sure that all the cryptographic storage techniques that are currently in use today that are best current practice need to be used. So I think that’s it. Scott, you want to go over the model itself?

SCOTT HOLLENBECK: Sure. Thank you, Andy. So, as the title of the slide says, the model that we've come up with as a proposal is based on two standards-based protocols, OAuth 2.0 and OpenID Connect.

Before I go through this slide full of technobabble though, I want to show you a picture. This little data flow-like diagram, it’s a bit of an evolution of the diagram that Steve showed you, with a little bit more detail about the interactions between the actors and the flows between the various data elements.

If you're familiar with single sign on services, the kinds of things where you go to access a web resources and you're prompted to sign in with your Twitter credential or a Gmail address or a Facebook ID, conceptually, you understand the model already. There's obviously a lot more detail than that, but the data flow is very similar.

So let’s now take a quick look back, and then we’ll come back to this. Alright, so there are a couple of prerequisites before these single sign on systems can work. First off, these additional service providers, they have to exist. There has to be some process to bring them into existence, and there's obviously software development work that has to happen for these services to be stood up and operational.

Requestors, which is the term that we borrowed from the EPDP to identify the people who are asking for data, must have credentials that are issued to them by an identity provider. And these identity providers are one of these new actors. Part of their responsibility when they issue these credentials is to associate identity attributes with the credential.

One of the nice things about this particular solution is that it works out of the box today using services provided by companies like Google and Microsoft and Yahoo who support OpenID and OAuth. Turns out though those providers know nothing about RDAP and so they have no association of these additional credentials yet. That's to come.

So once the prerequisites are met, the whole process kicks off when a requestor sends an RDAP request to an access service using some form of client application. The access service receives this request, and because the access service doesn't know who's asking, it sends a redirect to the client to interact with this thing called an identity provider.

The next thing the human will see is some sort of a webform operated by the entity provider where they're prompted to provide their credentials. Could be a username and a password, or it supports the use of client certificates if that's what the identity provider and client had negotiated ahead of time.

But let's just say for the sake of argument credentials are confirmed and validated. And then the next thing that the client or the human will see is a request to select various identity bits, these attributes that we talked about, and to provide their consent for this information to be shared with the underlying relying party, or the entity, the access service that controls access to this protected information.

So the requestor responds, fills out all these forms, pushes the submit button, and the identity provider returns something called an authorization code to the client and then sends another redirect, an http redirect to this thing called the access service, which starts the process of setting up an RDAP query.

The access service takes this authorization code and uses it to extract opaque blobs of data called tokens from the identity provider. The tokens are returned to the client. Now, it's these tokens that contain information about the identity associated with the requestor and some state information to determine authorization.

The client has the tokens, and then they send an RDAP query with this token information to ICANN's RDAP gateway. And when the gateway receives this information, it starts processing the actual RDAP query.

The gateway receives the query and the tokens, and then it sends both bits of information off to a third-party authorizer for verification. The authorizer processes these inputs, ensures that the identity information is valid, and that the matching of the query to the attributes is all good to go, and then returns a verification result to the gateway. This will typically be either, “Yes, good to go,” or, “No, that person’s not authorized for what they're asking for.”

So assuming that we’re authorized, the gateway will then send RDAP queries to the appropriate contracted parties RDAP servers, these being either registries or registrars as appropriate, to pull in all of the nonpublic data, the gateway processes and filters these responses to form a complete RDAP response, which is returned to the client, and then the client displays the result to the requester.

And again, here’s our picture in summary form. Same basic data flows. Passing on.

RAM MOHAN: Thanks. Scott, if you could just take a moment and point out the authentication provider, the authorization service, we've intentionally split them out, shown them separately, and I think it would be useful – we had some questions yesterday about whether the intent is to have them all be bundled together, whether they can be distributed, things like that. So I think it would be useful to speak a little bit about it.

in addition, we also had a bit of discussion in our deliberations about we got this idea of an identity provider and what that role is.

SCOTT HOLLENBECK: Sure, Ram, no problem. Yes, so those of you who are familiar with how WHOIS works today – and I assume that's pretty much everybody who's sitting in this room – will recognize at least two of the actors in this model. The client and the registry/registrar RDAP servers.

Well, that model doesn't work so well when you have to make decisions about identification, authentication and access control. And that's where these standards-based services come in. OpenID Connect and OAuth 2.0 are designed to give us the facilities that we need to properly identify clients, authenticate them and make access control decisions based on attributes associated by their identities.

But that means we need to add some additional players into the mix here, the first of which is this ICANN RDAP access service. We have, within our own little group, called this a proxy. So if you're familiar with how proxies work, you could certainly think of it in much the same way. It's a broker. It receives the queries and then decides who needs to be involved and how to vector the query appropriately.

But one of the first things the service has to do when it gets a query is it needs to know who it's talking to, and it does that through this authentication provider and this authorization service.

Now, the way the protocols work, these services can be performed by one entity who is sometimes described as an identity provider, or those functions can be split into different actors. The model that we described supports both methods of operation, and ultimately, it's probably going to be a matter of policy that determines exactly how this split is performed and which actor performs which function.

But as you can see the interactions there, the authentication provider receives the query from the RDAP service. It actually does interact with the client. This is the web-based interface that I described before. This is where the client provides their credentials. The authentication provider is the one that issued them, so it's able to perform the authentication function, and the RDAP service never has to be exposed to this information. It simply gets an attestation from the authentication provider as to whether or not the client is fully identified and authenticated.

But then that brings us to the authorization service. Once the RDAP access service knows that it's dealing with someone who is duly identified and authenticated, there needs to be a determination made about whether or not that requestor has the appropriate level of access to see what they're asking for. And as I said, that's traditionally a function in OAuth, that's performed [in an] identity provider, but it allows us to split that function into a third-party service. We're describing that possibility here in the model, and the way it works is that the query is sent, the comparisons are made according to some policies that are yet to be determined, and a thumbs up/thumbs down type of response is returned to the service which then acts accordingly in terms of building and querying the contracted parties RDAP servers. Enough detail, Ram?

RAM MOHAN: Yeah, that's very good. Thank you, Scott. One other thing, there was a question that has come through as to this ICANN RDAP access service that is there, whether in our model that means that the data that is sitting with the various sources, whether ICANN is going to get a copy of that data in our model or not. So it'll be good also to address that.

SCOTT HOLLENBECK: Sure. The data in this model stays with the authoritative sources. And I need to describe what authoritative means a little bit, because I know the thick WHOIS policy for example said that registries are an "authoritative" for the data.

We took a slightly different view here in that authoritative means that it’s the entity that has the relationship with the data subject. So it’s a matter of provenance, and where there data is most closely collected or where it’s produced.

So this is one of the reasons you're seeing a split and separation between registry and registrar functions here. Registrars will maintain the data that they are authoritative for. The ICANN RDAP service does not maintain copies of the data. The data transits through the service so that it can be processed, but there's no record kept other than through access logs. The data itself is not copied, it’s not maintained, it’s not cached. It’s transitory, and that’s as far as it goes.

RAM MOHAN: Thank you very much. I know you keep trying to push it my way. But I just want to stay a little bit on this because this is the core of what we were put together to do. One other question that came up, Scott, was this ICANN RDAP access service, the question that came up was in our model, do we imagine that that’s all centralized? Do we imagine that it’s just through a website, or whether we have other automated mechanisms? Number one. And number two, if you are an unauthenticated and perhaps even a not authorized request that comes in for public data or for restricted data, and it is data that is not a gTLD, what is our plan, our thought on that? So the bootstrap, the redirect, there might be some value there as well.

SCOTT HOLLENBECK: Sure. Okay, so the RDAP access service, we are envisioning it as a web interface, but with two faces. As Andy described, we have a need for online automated access, but there's also a need for asynchronous access, meaning you've got someone who doesn't necessarily have a credential, but they may actually have a legitimate purpose to request information. So there will be some sort of support for the client maybe filling out a webform and that form being reviewed and processed and a response being returned in some other way.

But as a web service, this RDAP access service can be implemented in any way that web services are typically implemented. Not necessarily one server. It can certainly be distributed in various places to deal with things like load balancing. It's really a matter of best practices for the support of http services.

The authentication provider and authorization service functions can be centralized, but they can also be distributed. And the model that we're kind of pitching here is one in which these functions are not centralized, that they are appropriately distributed. It makes a lot of sense for these functions to be performed by entities that have relationships with the requestors, because they know who the requestors are. For example, they're able to issue these credentials and make appropriate identification authentication decisions based on preexisting relationships.

And then with respect to the public data, you'll see that one thing we do have here is there is an expectation that the contracted parties will have public interfaces for public data so that clients will be able to send queries directly to registries and registrars. And what they will get back is whatever policy determines to be public data. Does that cover it?

RAM MOHAN: Yeah.

SCOTT HOLLENBECK: Okay. Thank you.

UNIDENTIFIED MALE: So if I may, in our document, we actually cover the different combinations of splitting out or combining identity providers and determiners of authorization, and we call that the actor models. And we have the set of combinations actually defined in the document. I believe there's four of them in there.

The other thing is going back to unauthenticated users or unauthorized users, in general, what we've asked for is that the ICANN RDAP gateway, when it gets a request for one of those, acts as more of a standard RDAP bootstrap server so that it can do an HTTP-level redirect to the source that is listed in the IANA bootstrap files.

One of the reasons we asked for this is not just from an ICANN gTLD perspective but RDAP is used in other contexts such as the RIRs and in the ccTLD space as well.

RAM MOHAN: Thank you. Thanks, Scott. So with that, I'll hand it over to you, Gavin, for the considerations segment.

GAVIN BROWN: Thank you, Ram. Yes, and Ram did mention this at the beginning, but I'll just repeat: as we were going through our deliberations and discussions, we identified some things that we felt had not necessarily been fully fleshed out prior to our beginning work. We also felt it out of our remit simply because we were closely focused on a technical solution and weren't very interested in getting involved in the policy.

So I'll just outline some of the items here. You can see there are a couple of slides. We'll go through them.

We've already briefly talked about data retention. As has been said, we don't envisage the access gateway having or storing registration data, in the jargon, it’s a reverse proxy but it’s not a caching reverse proxy. It doesn’t store anything, it just passes through the stuff that it gets from the contracted parties servers.

But there will be certain data elements that are stored. For example, the key one being logs. And we recognize that those logs may have some risk and value associated with them, and that it would be appropriate in the policy area for data retention rules to be applied to those logs and that that should be carried out.

Obviously, there are various things that the system needs to have in order to be able to for example ensure – we’ll talk a little bit on the next slide about transparency. It’s important to be able to audit the system to make sure that things are happening in good order. So therefore, logging is a key part of that auditability, but then there's a need to firstly reduce the risk of disclosure because the fact that someone’s submitted a request is potentially valuable information in and of itself, and so therefore having ap policy to reduce the risk of disclosure of that sort of information is highly appropriate.

It was also mentioned previously about service level agreements. Obviously, there are going to be a number of relying parties. The model we described has as number of potentially separate entities who each rely on each other’s services being available in order to fulfill their part of the system. And obviously, ultimately, there are the end users of the system, the requestors who are also relying on that system being available in order to fill their needs and fulfill their legitimate purposes for requesting the access.

So we've identified that a series of service level agreements should be defined and put in place to guarantee the availability and stability of the system.

We also recommend that ICANN Org should provide transparency on the performance of those service level agreements, and for example provide something like a status page where a requestor can see at a glance what the status of the system is.

We felt that it was important that ICANN should review what the potential impact would be of taking on the responsibility of running a system such as this, especially considering its role as a coordinating party, items three and four in this list can really be seen together.

There are obviously the legal risks, but also the operational risks, potentially significant scale. Again, depending on where on the spectrum the policy decision lies about how the system would be deployed, there are a number of other risks that need to be considered. Obviously, information security risks we outlined as well.

And it was important to us to flag for ICANN Org as well as the ICANN community what those should be and that certain reviews and assessments needed to be made to try and deal with those risks.

We did feel that it was necessary to flag up the issue about reducing liability to contracted parties. We're a bunch of geeks, we're not lawyers, so we can't answer the question about whether the system we're proposing does indeed reduce that liability. So we feel it necessary to encourage contracted parties to come to their own view, as I'm sure they will.

Mentioned briefly about transparency. We definitely feel that it would be a key part in ensuring the trust in the system that ICANN, if it chooses to run a system such as this, should be aggressive in being transparent about the way it's run and how it's used. Obviously, we're not saying specific requests should be published or disclosed, but statistical information about how the system is used, I think, is going to be a key datapoint in ensuring that there's trust in the integrity of the system, so we propose that a transparency report be produced on a regular basis for the benefit of the community.

And finally, we also recognize that there'll be outcomes from any kind of authorization process where the requestor may disagree with that outcome, and so therefore a mechanism for handling complaints should be put in place so complaints about processes or about system issues can be received, escalated accordingly and redressed through a complaints handling process, and that probably would also include deletion requests under article seven of the GDPR and other similar legislation where data subjects may either accurately or inaccurately come to ICANN and ask for data to be deleted about them. Ram.

RAM MOHAN: Thank you, Gavin. So if you could click me over to the next slide, unless you're back to self-service. Perfect. Thank you. So that actually more or less concludes our prepared comments for today's session. We are right now on track. We have community input that we're looking to gather from you now all the way through to Wednesday of this week, and actually, even beyond Wednesday but we'd really like to get your input. Our intention is to reflect upon that and incorporate that into the technical model. We plan to have several more calls as well as a face-to-face meeting in the middle of April to finalize our work, and we expect to publish the final technical model on the 23rd of April. And after we're done with that, this group will disband and we will have completed our task.

So, with that, it's time to move to questions from you. I see a gentleman there already. If there are other questions, please do line up. There are microphones on both sides, and I'll do my best to moderate. Sir.

RUBENS KUHL: I wonder whether the group considered a scenario where ICANN would only issue a token using OAuth and OpenID, and then the client would ask the contracted party directly? Because that would avoid data [flowing] through ICANN's systems. That would avoid most of the SLA issues. And even though it's not a caching proxy, a compromised proxy would still lead to data leak.

So there was the possibility of not putting the data flow centralized, but that was not chosen. Was there a reason for that?

ANDY NEWTON: Yes, we did discuss that. The reason we didn't go down that path was then you have to get into mechanisms for distributing policy to all the contracted parties, and you've put a much higher burden on the contracted parties to follow and to constantly update that policy. We felt it would be easier if ICANN.org were the place where all the policy filtering went on. So that's the reason why we went down that model, it's a simplified model from having access [inaudible] go directly to the contracted parties.

With regard to SLAs, I don’t believe– I don't know if anyone else does on the panel – that it actually changes any of the SLAs at all.

RUBENS KUHL: Even with a centralized policy engine at ICANN, the data flow could be not centralized. So there's no need for one decision to affect the other.

ANDY NEWTON: I understand what you're saying, but when it comes to SLAs, there's still going to be an SLA on ICANN, any services that ICANN’s running anyway. So I believe that the issues with SLAs do not change with either type of system.

GAVIN BROWN: If I may as well, I think the having ICANN as a single point of access does also address some concerns that we've received or we've heard from law enforcement about not wanting too much information about their own requests being passed to contracted parties. Maybe Benedict can talk a bit more about that particular point, or maybe Steve. But I do think when we were going through this discussion, that was something that I had in front and center of my point of view, that there has been a strong feedback from law enforcement that although they're fine on issues of transparency, there is a concern that – the fact that it’s a particular law enforcement organization or a particular officer at a law enforcement agency, that they're making a request and that being known to the registrar, that's a concern for them.

BENEDICT ADDIS: I'm going to be slightly more cautious and say the folks on the EPDP for phase two, this is a consideration that we are flagging right now for you, because there is a slider between pseudonymity of requests where all requests stop at ICANN and ICANN's responsible for logging, versus more disclosure, less disclosure over to contracted parties, and that is a policy discussion that law enforcement is going to need to have with contracted parties about anonymity of requests or pseudonymity of requests versus the ability of contracted parties to know who's asking, which may play into liability. But that is way above our pay grade, if that's an okay answer.

RAM MOHAN: Thank you.

BENEDICT ADDIS: There is one other benefit, which is that from a transparency point of view, having a centralized service allows you logging and the production of transparency reports, which I think we can all agree is a good thing.

RAM MOHAN: Thank you. The gentleman on my left.

KLAUS STOLL: Thank you very much. Klaus Stoll [inaudible] Group. A very quick comment and suggestion. Under considerations in point number six, transparency, you should maybe mention academic research, because a lot of people will be very much interested, and that goes beyond statistics. It would just be nice to have it mentioned. Thanks.

RAM MOHAN: Thank you very much. We'll take note of that and we will reflect upon that in our discussions.

VITTORIO BERTOLA: Hi. Vttorio Bertola from [inaudible]. I think my question was pretty similar to the one that was already made. I was trying to figure out why do you really want to have this centralized system [in the middle,] especially whether it's a technical choice, so you derived some of that technical decision that this would be better in a technical sense, or whether it is a policy choice, because in technical ways, if I had to build something like that, I think something decentralized would work much better in a lot of ways to just decentralize the [inaudible] and transmit so that they can be verified by the contracted parties, and we can discuss about it. I can even come up with proposals if you want.

But on the other hand, when the question was made, the two rezones I heard were policy reasons, like we want to keep track of everything for transparency, or law enforcement wants us to act as a proxy so that contracted parties don't get to see what they ask for. But these are policy decisions.

But in your presentation, you said that you cannot say whether these are actually better in terms of policy or legal responsibilities by the contracted parties. So now I'm a bit lost in understanding whether you're trying to solve a technical problem with the best technical solution or whether you have some policy requirements that make this centralized thing necessary, and in this case, what are they and can they be discussed?

GAVIN BROWN: So I agree with you, in some of those answers, you can kind of say there's some policy mixed in with that. But we did have good technical reasons, which was to lower the technical and operational complexity of the system.

Distributing policy and having the contracted parties have to constantly go update policy and understand what the policy language would be and how to parse it, however you would do that is a huge technical undertaking, and one of the reasons we did what we did was to try to reduce the complexity, the technical complexity in the system.

The other thing, and I didn't mention this before, is that – and we did put this in the document – with the way we've defined this, the contracted parties only have to use mutual TLS in order to know who to trust, whereas if we did access controls using a token that the client just hands directly to the contracted party, that raises the bar for what the contracted parties have to deal with.

RAM MOHAN: Thank you. Sir.

ALEX DEACON: Thanks. Yeah, I just had a question that's kind of related to Ruben's question, it's specific about system requirements 4 E and 4 F where it talks about passing attributes and identifier information onto contracted parties. And it seems to me that it's clearly stated that this is policy. I think you're going over the separation between policy and technology here.

But I'm curious, if the decision has been made that active model two is the one that we're going to go with, which I think is the right decision, and that model describes ICANN as the sole authorizer, it's not clear to me that we would need to send this data to the contracted parties. Can you maybe give me some context into some insight as to why these requirements – I'm not saying they're good or bad or disagreeing with them, just curious as to know the thought process as to why these requirements ended up [inaudible].

ANDY NEWTON: Yeah. Good question. The requirements are that the system must be able to do that, to have these identities and attributes of the requestor, and it must be able to support passing it through to the contracted parties if that's what policy dictates.

It wasn't really that that's the policy. So that's what we're trying to convey there. This is about a feature of the system that if policy says should be enforced or should be in place, then the system can't support that.

ALEX DEACON: Got it, so if the policy decides that it all happens in this ICANN proxy, then those attributes wouldn't have to be set down. The system would support it but we wouldn't use it.

ANDY NEWTON: Right, we need to be able to support it, and that's why we have the four different actor models. We're not sure which one is the correct one to do or if there are going to be multiple even. But that's why we laid them out the way we did.

ALEX DEACON: Okay. Thanks.

RAM MOHAN: Thank you.

NEAL MCPHERSON: Hi. Neal McPherson from 1&1 Ionos. I had a question with regards to the historical data. I think Steve mentioned that the whole process doesn't have to happen in real time, there will be use cases where there could be a need to go out and get claims or get information or whatever, it could take a long time. So, is there a timestamp or anything like that that you have to deliver the data based on today or yesterday or when the claim first came in? And also with regards to requests, a lot of requests that we get are based on, hey, who is the domain owner [inaudible].

RAM MOHAN: Could I ask you to get a little bit closer to the microphone? I heard use cases, I heard historical data, and then I'm filling in the blanks and I'd rather not.

NEAL MCPHERSON: Alright. Yeah, that's much better. So with regards to timestamps, yeah, the process can be out of [band,] it can take a long time. What timestamp of the WHOIS data or RDAP data do you have to give out based on a process that isn't happening in real time? And also, I was saying that we get a lot of requests for historical data. How is that worked into the process, that a requestor says, hey, I need the data from six months ago, who was the domain owner?

ANDY NEWTON: So, I work for ARIN, and we have this thing called WHOWAS, which is basically, "Give me what the registration data looks like six months ago" or whatever. We did discuss things like bulk WHOIS, WHOWAS services and so forth, and we kind of ruled them as out of scope, at least for now, because that basically complicates to a degree which we're not sure is within what was in our remit.

But we did discuss those things, and we said, no, let's set those aside for right now. Does that answer your question?

NEAL MCPHERSON: [Thanks. Yeah.]

RAM MOHAN: Next question, please.

GREGORY MOUNIER: Hi, I'm Greg Mounier from Europol, and I want to ask a question which his related to this one, it's about another feature that criminal investigators are using all the time, 80% of their research, and that's the reverse searches. So for those who are not familiar with this, it's just to be able to cross-reference or to identify all domains that have been registered with one specific type of information. Could be an address or the name of a registrant.

So I've read in your reports that technically, this was possible to develop, but for some reason which I haven't understood, the TSG hasn't said to develop it. So my first question is, what would it take for you to decide to include that in the scope of your study? And then the reason why it is not included in the study. Thank you.

ANDY NEWTON: Yeah, so the basic reason why it's not included is because reverse search is not part of RDAP at the moment. There is a draft in the

IETF which will actually be discussed in Prague in two weeks dealing with reverse search, and this was just never part of the base RDAP to begin with.

Beyond what that draft says and how it would be supported, there are questions about how you go about getting that data from all the contracted parties in the space. Did you want to say something, Benedict? So yeah, that's the reason why it's not covered.

GREGORY MOUNIER: If I may just, I think it would be a really big pity if at the end, the policy were to say, yes, you can do it, and on technical [side,] we will say, "Oh, we haven't dealt with it for some reason. So it would be great if at the end of the process, we could actually do it technically.

RAM MOHAN: Thank you. We're going to get to you in just a moment, please stay at the microphone. It's 2:46. On 11th March 2011 at 2:46 PM local time, a 9.1 magnitude earthquake struck in the Pacific Ocean of the northwest coast of Japan's Honshu island.

The earthquake known as The Great East Japan Earthquake triggered a massive tsunami with waves that rose to heights of up to 40 meters and traveled up to ten kilometers inland.

This was the most powerful earthquake ever recorded in Japan and the fourth most powerful earthquake in the world.

An estimated 20,000 people were lost, and close to 500,000 people were forced to evacuate. In remembrance of the lives lost and affected by the The Great East Japan Earthquake, we will now observe a moment of silence.

Thank you. Andy, would you like to respond to the follow-up?

ANDY NEWTON: I'm sorry, I lost track. What was the follow-up question?

GREGORY MOUNIER: It wasn't a follow-up question, it was more just a statement saying that it would be a pity if technically it wasn't possible if the policy would say it was possible.

ANDY NEWTON: Hopefully, future features that are in RDAP can be used with this. Yes. That may be a matter of policy. There may be some more technical things to work on, but yeah, that would be a good thing to support in the future if the community really decides they need it.

RAM MOHAN: Thank you.

SÉVERINE WATERBLEY: Hello. Good afternoon. I'm Séverine Waterbley from Belgium. I'm a GAC member. If I do understand correctly, the ICANN will be the processor in the meaning of the GDPR, and the authorization services will be the processor [of] the processor. So are there contractual relationships foreseen between the registries and the services to provide us authorization or authentication? Thank you.

RAM MOHAN: I think that's a very good question, and I don't think we're qualified to give you an answer on that. But what we will do, just to make sure that this isn't lost, we will record it, we will make sure that this is passed on into the folks at ICANN Org, because the legal aspects of it, etc. are not things that we have devoted any of our time or a great deal of our time and energy on.

BENEDICT ADDIS: [And it's subject to a certain amount of discussion in the EPDP.] And again, just speaking with my EPDP hat on, there has been a great deal of discussion there about the exact nature of the legal agreement within ICANN and the contracted parties. And I know that ICANN Legal have opined on that issue. So I think EPDP, and particularly phase two, which convenient enough, there will be some discussions in this very room in 40 minutes.

TIM CHEN: Okay. Thank you. Tim Chen from DomainTools. First of all, thank you for the work that you've done here. I think it's a good service to consider the technical realities of getting this done in parallel with the policy development, so I applaud you for your work [like all] volunteers. So thanks for that.

Just two quick technical questions. One is there was a note at the bottom of the early slide on considerations that mentioned – I believe the term was multi-use queries. I'm just curious what that means. It was early [– Ram, it might have been your section,] or the one right after. It was a long list of like 13. There it is, multi-use requests. Sorry.

Can you just help me understand what that is, the bottom of that slide?

RAM MOHAN: Sure. So this is in the early stages of our work. What we were looking at was that the request for data could take many different forms. One form might be that it's a party that is authorized to only access data for one data element once, and in other cases, it might be that the authorization is more persistent, but also, it's persistent but for one class of requests or to be able to only access a certain number of data elements.

So, it was not clear to us at this part of our deliberations that we should restrict the model to it must only handle a request once per data element and once per object, or in either a persistent way or in an ephemeral way. So that's really what the multi-use request questions are about.

TIM CHEN: Okay. Thank you. Second question, if I may. The bootstrap service was mentioned yesterday in the sessions that you did with the contracted parties and then today. [I'm not technical,] I did attempt to look at the RFC for that [in a few minutes here,] but I think we clarified with the greatest question that that's not somehow reaching – I think the term yesterday was reaching across contracted parties. It's not doing a reverse service. It seems like the basic point of a bootstrap service is to find out which authoritative source to go to. But getting back to the comments that were made briefly about bootstrap – it might have been Scott that said this – was there [an application] that this service will handle requests for information outside of just gTLD domain name data? And if so, could you expand on that?

SCOTT HOLLENBECK: Sure. In theory, it's possible. I know we jokingly talked about, "Wouldn't it be great if internic.net actually performed some useful services again these days, right? So no, indeed, how this ends up evolving over time is going to be a matter of policy coordination and implementation. But in theory, yes, that's completely possible.

TIM CHEN: Thank you.

ANDY NEWTON: I want to add to that. There is a fairly good likelihood that there are people who will write RDAP clients that they don't want to do the bootstrapping process themselves. They've probably done this with [Kerl] and Bash or something like that, and they're just going to pick a bootstrapping source and it might be ICANN. In that case, it would be great if ICANN acted as a typical bootstrap server using the IANA files.

RAM MOHAN: Thank you. Yes, please.

UNIDENTIFIED FEMALE: Just an observation. I think the only party really that was excluded from this model and from the EPDP was the end user. I know that you have a use case where you can check your own data, but this model, I feel that the end user is a canary in the mine, and we're always sacrificed. I'm just an end user, I don't belong to any group, and it is unfortunate. I feel that this is something that we have to fight for, but we are always excluded, we always have the worst deal, and all the parties will find mechanism to either get the data they want with restrictions, through legal means, through authorizations, but we're the only one who has been excluded.

It is unfortunate. I feel this is something that we will have to fight for, and it's just the beginning. Thank you.

RAM MOHAN: Thank you for that. Kind of a peculiar issue to bring in front of us, because in all of our discussions – and if you'll go back and look, listen to the recordings and transcripts and things like that, you'll actually find that we devoted a significant amount of time looking at it from the point of view of the end user who is making a request. And although the use cases, when they boil down, we have five use cases, but there was actually quite a lot of discussion about the end user, and in fact the principle that Andy was talking about, keep it simple, and try to not have multiple services and have the end user go to multiple places to get information. That was directly focused on making sure that the user journey and the user experience is something that is kept in the forefront.

So, at least speaking from my point of view, it feels like the end user has been a consideration for us. But certainly in the overall process, I agree that all of us have to look at what are the needs and requirements of end users. Are there others? Steve, you have your hand up.

STEVE CROCKER: Yeah. I'd like to drill down into your question a little bit. What is the actual issue here? Presumably a registrant can go into the account with the registrar and find all of the information and verify them. And they have direct access there.

So I'm puzzled, genuinely, as to what the issue is that we're talking about here that is meaningful in terms of supporting the end user.

RAM MOHAN: And Benedict just helped me, he said perhaps what you're talking about is the data subject rather than the actual end user.

UNIDENTIFIED FEMALE: Yes. Exactly.

RAM MOHAN: And that is use case five in there that we've talked about.

UNIDENTIFIED FEMALE: May I add that I did not know about WHOIS before that? Now that I know WHOIS exists, I would have used it. I would have wanted to use it. The lady from the GAC, French lady who said that she was trying to push for consumers to have access to information for companies is one thing. That is very important for us.

And I think the solution would have been actually for us, if we were the center of this, would be to develop the use of WHOIS for us as a protection and not to eliminate it. Yes.

BENEDICT ADDIS: You know, after our initial confusion over terms – and I'm going to speak not as a member of the technical group but as a member of the EPDP again, I think you make an excellent point. And I'd be sad to see such a system become a boy's club for a small number of people that have jumped through the [regulatory] hoops.

And I think that the joy of this model is that it does have the flexibility, and again, we're going to sue that word again, if the community decides that it be used in such a way for ordinary Internet users, to find out who they're transacting with. That is a policy that is something that you must go to the community with.

UNIDENTIFIED FEMALE: Yes. Right. I'll fight for it.

BENEDICT ADDIS: Thank you.

RAM MOHAN: Thank you. You will have the last question.

UNIDENTIFIED MALE: Thanks. I'll be real quick. And it’s kind of related, and it was as comment that I had on the user journey section, section 4.1 of your report. It describes a user journey that I'm not familiar with, and I've been using WHOIS RDS for a long time, specifically the second bullet. I've read that several times, and I think there are some improvements that could be made to that user journey to really put it in the context of a user of WHOIS, or RDS, or whatever we’re going to call it moving forward, and just make it really clear as to what user we’re talking about and what journey they're taking, and really put yourself in the shoes of the person that’s going to be using the service to look up RDS data.

RAM MOHAN: Thank you. That’s terrific feedback. Appreciate it. We are almost at the end of this session. What I would like to do is to pass the microphone to the various members of the Technical Study Group for just a moment if you have any other things that you'd like to say before we wrap the session up. I'll start with you, Scott.

SCOTT HOLLENBECK: Sure. Thank you, folks. Remember, this is a consultation. We threw this out there. We think it works, but we are certainly looking for your input. Please send us your comments, we want to hear what you have to say.

RAM MOHAN: Andy?

ANDY NEWTON: I just want to echo what Scott just said. We're looking forward to your comments, just please send them in and we will discuss them. Thank you.

RAM MOHAN: Benedict?

BENEDICT ADDIS: What a team. Thanks, everyone, and thanks for coming to listen to us.

RAM MOHAN: Jorge?

JORGE CANO: I want to say thank you for your comments. We really appreciate them.

RAM MOHAN: Jody?

JODY KOLKER: I don’t have anything to add. [That’s great.]

RAM MOHAN: Gavin.

GAVIN BROWN: I do have one thing I’d like to say, which is based on previous comments about use of registration data as consumer protection, enabling consumer protection. I think one of the amazing things about RDAP is how much more accessible registration data becomes when it’s provided over RDAP rather than over Port 43 WHOIS. If you're writing a web application in JavaScript or you're writing a mobile app for a phone, you have almost no way of getting access to Port 43 data unless you're using some sort of proxy server that you have to stand up and run for your users forever. RDAP runs on the web and the data is provided as JSON. Any programmer knows what JSON is.

I think that the adoption – and this [is kind of sideways from what] TSG is doing but speaks to RDAP as it's being adopted. I think it can have the potential to enable a great deal of benefits for consumers who want to obtain trust in any identifiers that they're using, because it will allow the key information about those domain names and other resources that RDAP makes available to be available in much more easily to the end user. I envision things like a browser extension where you can click on a little icon in the address bar, and the browser can natively load this data from the registry or registrar RDAP service and display in the browser without having to go through a clumsy Port 43 system. And you can use – the benefits of the web framework allows us to have things like caching and security to give you a great deal of integrity in that system.

RAM MOHAN: Thank you, Gavin. Tomofumi.

TOMOFUMI OKUBO: Thank you very much for staying awake until the end of the session. I'm looking forward to –

UNIDENTIFIED MALE: Not everyone did.

RAM MOHAN: Steve.

STEVE CROCKER: Ditto on everything. Thank you. Diana? [Eliza,] Gustavo? John.

JOHN CRAIN: Just one thing, amazing thanks to this team. As engineers, we always hate having to build systems without requirements, and that's precisely what this team has had to do. And I think they've done a fantastic job, and these folks who are, like all of you, volunteers, There's a lot of heavy lifting to build something that we hope was – the term I use is "policy agnostic." It's not a very good term, but it's flexible enough that it will survive the beatings it will get over the next years, and a lot of heavy lifting in a very short time. Amazing work, guys.

RAM MOHAN: Thank you. That concludes this session. Appreciate all of you for coming. Thank you.

[END OF TRANSCRIPTION]