ABU DHABI – ICANN GDD: RDAP Pilot
Wednesday, November 1, 2017 – 13:30 to 14:45 GST
ICANN60 | Abu Dhabi, United Arab Emirates

MARK ANDERSON:	Just to make sure everybody is in the right room, this is the RDAP Pilot discussion group. Everybody in the right room; this is what you're here for? I'm seeing nods, that's a good sign that I'm in the right room, because that would be embarrassing. First I'd ask for an indulgence, this is the second meeting of this group. We're a fairly new group, we're just getting our feet underneath us so if it seems like we're a little disorganized, that might be because we are. But as my friend Kal said, this should be very exciting because we're solving all the problems of WHOIS here, so very exciting group to be in.

I asked Francisco to start things off for us with a little bit of background, and if anybody was at Tech Day, I apologize, because you're going to see the same presentation, but Francisco gave a Tech Day presentation on RDAP; gave a little overview on how we got to where we are so I thought that would be a good way to start off this meeting, provide a little background for everybody, and I see a lot of new faces, or faces I don't know so if you're not familiar with RDAP and what the RDAP Pilot is about, this will be a good way to start things off. So I'll turn things over to Francisco.

FRANCISCO ARIAS: Thank you, Mark. Hello everyone, this is Francisco Arias from ICANN org. So this is just a brief overview of the end of the session, and I'm going to just do the background introduction about RDAP. So a brief history of RDAP; this started in 2011 with the Security and Stability Advisory Committee published SAC 51 in which they called the ICANN community to evaluate and adopt a replacement protocol for WHOIS.

In the same year, 2011, the board adopted that recommendation and instructed the staff to work on a road map to implement that recommendation, which was published in 2012. And in the same year, work started in the ITF to develop this new protocol that will replace WHOIS; that eventually was named as RDAP. By March 2015, the protocol was finalized in IDF and in that same year work started in the context of ICANN to develop what we call gTLD RDAP Profile; the profile being the set of functionality from the RDAP protocol that was going to be implemented in the gTLD space.

For those of you that are not familiar with RDAP, you can think of it as a menu, defines how to do things, certain functionality, but it doesn't tell you which ones you turn on, and which ones you don't. And so the profile is the one that tells in this case the

gTLD space which functionalities should be turned on, and which ones should not be.

So we published that on July 2016, and unfortunately a few days later we in ICANN received a request for reconsideration regarding the inclusion of those RDAP requirements in the consistent level and display policy among other things. So in response to that a new version of that policy was published, removing the requirement for RDAP and for a period of one, ICANN org. worked with the [inaudible] group and the registrars trying to find a way to move forward with RDAP implementation.

So that came finally in a form of a proposal from the [inaudible] group with support from the Registrars Stakeholder Group that was sent to ICANN last August, and a month later ICANN replied accepting that proposal.

So let's see now what is included in that proposal; so the first element calls for gTLD registries and registrars to work with ICANN on developing a profile or profiles, so here the concept is being introduced, the possibility of having more than one profile for the gTLD space, and also to develop a timeline to implement RDAP in production. The proposal also called for having a pilot program that would run until July next year, and at the end of that period to have an established target to move to production.

ICANN 60
ANNUAL GENERAL
ABU DHABI
28 October–3 November 2017

Some of the other elements were a little bit in more detail, called for ICANN to provide a waiver for a sector requirement that the registries have in their agreements in regards to new services. Another important point was that the intention for the pilot was to give freedom to the contractor parties to implement whatever functionality they deem appropriate, as long as they comply with the RDAP RFCs.

There was also encouragement to the contractor parties to publish the terms and implementation features that they are using, so that users can know what to expect in a given implementation. And there are a few other elements, I'm not going to go into details on those; the presentation is available on the website if people are interested.

So at high level, this is the timeline; on 5 September, so a few days after we in the ICANN org. accepted the proposal, we started the pilot as requested in the proposal, and remember this pilot is intended to run until the end of July next year. At the same time the work on developing that profile or profiles also started as Mark mentioned, and as a reminder, the intention is for this pilot to end by July next year. By then we should have a profile, or set up profiles to be used in the gTLD space and there should be an agreed timeline for moving RDAP to production. The date when RDAP goes into production is still to be defined.

For those of you interested in participating, we have a pilot webpage; you can see the link there. In there you can find more information of how to participate in the pilot. It's important to mention, it's open to 21 contractor parties to participate offering services and for users to participate using the different implementations. As of yesterday last time I checked there were five contractor parties participating in the pilot. Those that were four registries for a total of 50 TLDs, and we have one registrar listed there. Hopefully we will get more that are interested in participating.

There is also an ICANN RDAP page that contains more than just information about the pilot, and this is the session that we are in here. And that's it, thank you.

MARK ANDERSON: Good, thank you, Francisco. Again, for those people who came in since we started, my name is Mark Anderson from VeriSign and I agreed to facilitate this session today. I do want to highlight a couple of the things that Francisco mentioned in there; I think the SAC 51 report is really critical in how we got here, it was one of the key pieces of documentation that highlighted the shortcomings of WHOIS, and I think it's generally understood within the community that there are shortcomings of WHOIS, and the SAC 51 report is really the report that

triggered the development of a replacement protocol. That replacement protocol is RDAP.

In that report, it called for registries, registrars and ICANN to work together to determine how to implement RDAP. So this group here is attempting to do just that; we're looking to figure out how to implement RDAP hopefully in a way that addresses the shortcomings of WHOIS.

I think it's also important to note that this is not the policy making process; that is occurring as part of the Next-Gen RDS PDP; that policy making process actually meets later today at 4, if you're really interesting in getting more about it, but that is the process that will develop policy around how to implement a next generation RDS.

Here we're working on the technical, or the technical implementations of a pilot, and of an implementation and sort of following through on the SAC 51 mandate for registries, registrars, and ICANN to work together on that implementation.

As Francisco mentioned, we have five pilot participants so far and hopefully we get more people interested and participating in that. But what's not on the sign-up page is end users, and so hopefully many people here are interested in participating in the

pilot as an end user in providing feedback and input into the development of RDAP.

And that's part of what we're doing here today, is giving the broader community a chance to hear what we're doing and provide input into what an RDAP implementation should look like. So we'll have a chance a little bit later to provide some input on what should and shouldn't be in an RDAP pilot, and I look forward to hearing from many of you.

So I'd like to take a moment, I know we have a couple of registries represented here who have pilot implementation so far. My company is one of them; we have an RDAP pilot running for both .com and .net. The information is available via the links that Francisco provided, you can figure out how to get to our RDAP pilot implementation.

At this phase though, it's a very vanilla implementation. We've implemented very base functionality for RDAP and we're looking at that for sort of a starting point for RDAP. We don't anticipate that to be the final version; it's a place to start. And we have during this pilot period until July 2018 to try things out, see what works and what doesn't work, and give the community a chance to participate and provide feedback in that.

Since Stephanie is in my direct line of sight, I'd like to pick on her and ask her to tell us a little bit about the Google implementation of the RDAP pilot.

STEPHANIE DUCHESNEAU: Sure, and my colleague, Brian Mountford, who actually built our implementation is on remote I think -- do we have remote enabled for this? Perfect, so I'll speak quickly to it and then probably turn over to Brian who can probably provide a little bit more technical insight on it.

So we actually built our RDAP deployment pretty closely to the ICANN operational profiles, so some of the concepts and the rules that you'll see in there sort of dovetail very closely with what we've built, but in the month or so leading up to actually deploying it, and we just did it a few days ago, we made a few changes to how we plan to implement in light of some of the conversations we'd been having around GDPR.

We haven't committed to whether or how we would use RDAP as part of a GDPR solution, but for trial we're now, for all users querying they're going to see a redacted input, and Brian can correct me if I'm wrong, but I believe our output for all users querying is basically what you would see in a thin registry, but

we're also including the registrant tech and admin ID's, but not the other contact information that's associated with them.

But for registrars that are logged into our registrar console, you would be able to say the full WHOIS information for your domains under management. So a registrar would be able to see it for their own domains, but not for domains that were managed by other registrars. We're still running WHOIS in parallel, but sort of just doing an initial test for how a gated system might work for something that we were already credentialing.

I think, I won't speak too much unless Brian has more stuff he wants to add, I also note that like our open-source registry platform, Nomulus, it's open-source, so there's a lot of documentation about how it's running in our GitHub repository, which we've linked to in our application to ICANN. So I'd point people there if they're interested in knowing more, and turn it over to Brian if he has anything he wants to add. [AUDIO BREAK]

MARK ANDERSON: Brian, are you with us? Okay, we'll give Brian a chance to respond. Do we have a question down on this end of the room?

JIM GALVAN: Not a question, I just wanted to get on your list Mark.

MARK ANDERSON: Thank you, Jim. I'll give Brian one more call. Do you have a question over there?

JIM GALVAN: Oh, you do have a question down here, yes.

MARK ANDERSON: Go ahead, please.

JONATHAN MAKOWSKI: Makowski; Risk IQ. I was curious whether or not you're going to be exploring technically within RDAP how authenticated access would work. The end users; would we have an opportunity to test out some solutions with that?

MAXIM ALZOBA: I have two related questions. First of all, RDAP, tiered access, and have you thought about the solution for authentication of law enforcement, and the second question is tied to this; I suggest conducting browser form, so people can something more readable I would say. Because simple policemen, they are

not good at reading text, and parsing it in their head, and because as registry and on our registrar site we spend literally years teaching them how to use Whois.

And if we can with solution, which is not compatible with law enforcement, they have a bad habit of making you compliant using their methods, so it's a good idea to start development. I think it's technical, because browser form, the methods might allow you to read the output of RDAP, of JSON in some human readable form I would say. And it's important for all of us. Thanks.

STEPHANIE DUCHESNEAU: This is Stephanie Duchesneau with Google. To respond to the prior question, and I'm sorry, I don't know your name, I would regard exploring how gated access could work as part of this RDAP pilot group, and something that we're exploring as part of this implementation, but I think the question of how credentialing processes for different groups would work, and in the context of Whois, who should have access to what data, I would regard those as being policy questions or legal compliance questions in the context of GDPR that are out of scope.

So there's like two different sides of it; on the one hand how it would work technically, which I think is something that we need to be looking at but the question of who gets what and how do we figure out who is whom, I don't think that that's within the group's remit.

JONATHAN MAKOWSKI: Jonathan Makowski; Risk IQ. I agree with that; I think we have to figure out a way to move forward with technically exploring all aspect of RDAP while the policy issues are being worked on. I just wonder how to do that.

STEPHANIE DUCHESNEAU: Yeah, I totally agree and that's why we've chosen to do it with something that's fairly non-controversial and we're already able to credential.

JIM GALVAN: So Mark, can I jump in here to talk about our implementation? Yes? Okay. So, Jim Galvin from Afilias, and listening to this discussion here about authentication, and I prefer to use the phrase "differentiated access," not "gated access," so I get kind of fussy; it's my little pet peeve about just trying to make sure we use terminology that applies kind of broadly.

But we actually have a prototype implementation that's running now; it's not our production [inaudible], it's fairly simple; it's running off the .info TLD, and you can find a reference to it on the sheet of participants if you go to the community page, there's actually a URL to the help file for the RDAP server, so if you just click on that it takes you right to it and then it explains how to add other commands to do stuff there, so all of that actually works.

Our primary reason for getting involved in this; we were actively engaged in the development to the RDAP standard that the protocol itself, so we had gotten that far. The thing that we really want to get to in the pilot, what I hope to get out of it is credential management, and dealing in face with differentiated access.

We see that as the future; this is obviously something that's going to come, and it's going to have to be there, so I'll phrase it in the following way; I mean our goal is actually to demonstrate that you can have families of credentials, and therefore you can have a different profile of response for each family of credentials, and show that all of that works, and works across participants. So we'll create a pilot that allows us to demonstrate that all those features work.

And then as Stephanie said, I absolutely agree that who gets a credential based on what criteria and what that profile actually looks like; those are all policy questions. We'll experiment with some options and make sure that we can make things work in a variety of different ways, but I expect the next generation RDS PDP working group to be responsible for at least that working group or something like it along the way.

I guess I don't want to presuppose where all things are going to come from around here, but there will be some policies that tell us what's going to happen there. So law enforcement will get covered in some way, probably other families of kinds of things, we'll just see. So, thank you.

MARK ANDERSON: Thank you, Jim. Maxim, do you have a comment?

MAXIM ALZOBA: I'd like to rephrase my question in a more technical arranged way; have you thought about authorization across different registries for some families of those I'd say users with differentiated access?

JIM GALVAN:	So I'll tell you my thoughts about where we're going to go with our implementation; I think that the larger context for your question is that that's really a policy decision as to what that stuff is going to look like. What I plan to do, because we host almost 200 TLDs, so we already know we're going to have a variety of different policies that are going to come to bear on us; GDPR is sort of the over-arching one that everybody things about, it's the latest buzzword this week. We'll see what's coming up next, but we fully expect a variety of different kinds of profiles, and we're going to have to build in the flexibility to be able to set those things according to whatever policy is driving us.
	I fully expect a variety of Sovereign Nations if you will are going to have their own ideas about what's allowed to be displayed or not displayed. And who knows what other criteria are going to come at us, so we're just planning on some flexibility and choosing that flexibility is the vision that we all have to just sort of choose, as we watch the rest of the policy stuff develop around us.
MARK ANDERSON:	Go ahead.

OWEN DELONG: I think one of the things that's technical in that respect though is I think we need to consider the ability to credential across registries. And by that I don't mean across the different registries that Afilias operates; I mean a way to have credentials where a group of users can be credentialed and in such a way that it affects VeriSign and Afilias and all of the other registry operators, some sort of shared credentialing across that I think is a desirable possible technical thing.

JIM GALVAN: So let me phrase that in the following way; to my mind, that's covered by the idea of ensuring that the implementation can work with families of credentials. And the idea is that for any given family, all registries and potentially all registrars will validate the credentials in that family. So the existence of a family and who gets a given credential within a family, that's a policy issue and that'll be settled elsewhere.

I mean, we're going to base this on certificates, as far as I know, no one has suggested anything different. It's sort of the technology of the day and so that's what we're going to do. And so you can imagine that these families are just going to be independent PKIs and we all are just, anybody who's going to have an RDAP server is just obligated to just validate that particular certificate, and to understand how to do that.

I mean the model here to think about is what browsers do today, in order to show you that cute little lock or different color that's in the address bar and stuff, they maintain a cache of root certificate authority certificates, and I imagine that RDAP servers are going to have to do the same kind of thing. And so all of that technical system will have to come into existence too, but the driving force behind all of that will be policies, so I want to build an implementation that will handle all of that, it's all a known quantity if you think about how browsers work and that kind of thing, and then whatever the policies decide we'll be ready for it.

MARK ANDERSON: Thank you Jim. We have a couple of people that want to get in the queue; sorry I don't know your names.

GREGORY MOUNIER: Hi, I'm Greg Mounier from EUROPOL; just a practical question on the pilot project. So, as far as I understand, there's no one portal where as an end user if I want to participate to this program I would have access to all the registries, implementation projects; you have to go to each registry and test their prototype, yeah?

JIM GALVAN: Yeah, that's correct. The pilot itself is voluntary; it's open to any registries and registrars that want to participate, but by virtue of that it's each registry or registrar's individual implementation.

GREGORY MOUNIER: But there's no plan to actually make it more user friendly so that we have one place to go and then we can test all the profiles?

MARK ANDERSON: I'd say at this point, no there's no plan for that. I think the gentleman next to you had his hand up first, so I'll come back down there.

IRANGA KAHANGAMA: Thanks. Hello, this is Iranga Kahangama with the FBI, yeah from a technical perspective I kind of just wanted to echo Jim Galvan's thought that building in those capabilities would be really interesting to see from a law enforcement perspective, and then obviously switching to the policy track to figure out all the specific details.

                         But that also from a law enforcement perspective we would be interested in just getting some of our analysts and agents as end users to kind of dabble in, and then report any feedback as needed.

MARK ANDERSON: That's great to hear, thank you. I don't want to set any expectations that we're going to get it perfect out the gate, but the more participation we get, the more feedback we get, the more opportunities we have to work out some of the implementation challenges, so it's great to hear and we look forward to your participation and input.

IRANGA KAHANGAMA: Yeah, and to the extent that it's helpful, we could talk maybe offline about implementation challenges that we would foresee that may be worth considering so that you have it up front for whatever technical discussions you're having.

MARK ANDERSON: Great, thank you.

JONATHAN MAKOWSKI: Jonathan Makowski; Risk IQ. I wanted to throw out there for consideration that what we might find as a community now is that the existing Whois protocol is going to adapt to GDPR in certain ways that move faster than the RDS working group does with this next generation that RDAP is supposed to sort of accommodate, so it's possible that we really need to explore

technically, or be prepared to explore technically, adapting RDAP to the GDPR issues that get solved more quickly than the RDS working group is able to solve; those policy issues.

So we would welcome working with the community on making sure that technically speaking, as those problems get solved in the near future, that we can actually use the protocol technically to try to test out what was discussed just earlier with how this is going to work, instead of waiting for the next generation RDS to come up with the policy issues.

Because it may be that the Whois protocol doesn't end up from a policy perspective needing to be changed in a way that's contemplated, and if that happens, then we need to basically test out technically how we're going to make that happen in RDAP.

MARK ANDERSON: Thank you, those are real good points and it's almost like you're reading ahead on the slides, so that's great to get that input. We have another couple comments down there?

DAVID PEALL: David Peall from DNS Africa, Registry Service Provider to a couple of registry operators, one being RyC in Europe and we're

also developing our RDAP service at the moment, we're on the list that Francisco provided earlier, and we'll be publishing our timeline on the URL provided. As with everyone else, we're building a dynamic system that will be able to load profiles depending on the region that the information is served from.

And while we might be talking a lot about policy in these meetings, the legal constraints are going to potentially overtake the policy development, and we will obviously have to consider our risks in terms of breaking local laws, versus having trouble with ICANN, and I see that sentiment has come out a few times.

So with regard to that, we are developing profiles to obey both the European laws as well as South African laws because we have customers in those region as well; very similar issues in terms of protection of private information.

MARK ANDERSON: Thank you. I know there's another comment down there, I'm going to ask you to hold for a second, I believe we have some comments in chat. Antonietta, can I ask you to read them?

ANTONIETTA MANGIACOTTI: Yes, we do. We have one from a remote participate; Brian Mountford. His question is, are we considering including some

form of the bootstrap functionality in the pilot program that would be a cross provider facility related to what the speaker was requesting?

MARK ANDERSON: I'll jump in on that. That is actually possible; there's an IANA bootstrap service that exists for RDAP today. The VeriSign pilot is actually registered in that IANA bootstrap service, so that's a good point. That option is available today, and to your point, I'm sorry I forgot your name already. To your point; that would facilitate the creation of a cross pilot functional client.

ANTONIETTA MANGIACOTTI: We also have another question from Houssem Kaabi; does the deployment of the RDAP mean the drop of Whois?

MARK ANDERSON: Unless somebody else wants to jump in, I'll go ahead and tackle that as well; no, it doesn't. And deliberately so, I think the general feeling is that that is a question that belongs in a next-gen RDS PDP. There's certainly a lot going on in the RDS space today, so I don't want to put a crystal ball in front of me, but the charter of the next-gen RDS PDP does include a question on what the sunset or ultimate replacement of Whois should look

like, and so at least at this time I think that's a question best left to that group.

JONATHAN MAKOWSKI: Jonathan Makowski; Risk IQ. I would just clarify that if part of that charter also to adjust whether it even needs to be changed and how, so the assumption is not that Whois is dead and we need to replace it with policies applying to RDAP. They need to work through the policies to decide whether that's true in the first place, and I think the community is going to move a lot quicker than that group.

So I would say that we should just make sure that we test out the possibilities of the protocol as quickly as we can, being flexible and open-minded to how to go about doing that phase.

MARK ANDERSON: Very good point, thank you. And just a note; one of the stipulations for the RDAP pilot is that when and if the next-gen RDS PDP finishes its work, any pilot implementations would adjust to the outcome of that. So certainly the first question as you already pointed out of the Next-Gen RDS PDP is to answer if a replacement protocol is even needed, but again I won't put my crystal ball in front of me on that one. Go ahead.

DAVID PEALL: Hi, David Peele here again from DNS Africa. I might be coming in a bit late here in terms of how the RDAP servers are identified through a central information area, and you mentioned something earlier. But is there any possibility we can use something like an SRV record under a demand that would easily identify where the RDAP services are available for the TLD in question? It makes sense to me.

MARK ANDERSON: I might have to ask Jim Galvan to bail me out on that question.

JIM GALVAN: Can you repeat the question?

DAVID PEALL: Using something like an SRV record in the DNS to identify where RDAP services are to be found for the TLD in question?

JIM GALVAN: Yeah, bootstrapping is an issue. I mean as Mark had said earlier, right now there's an IANA registry that does provide currently the way to bootsrtrap access. It is one of the deficiencies overall in the RDAP protocol. I think I said before, the credential

management is the big thing that we're interested in, but there are some other sort of lesser technical features that are important to understand and work through.

One of them is the bootstrapping, another one is redirection. That's something that will also have to be supported and worked out because it's not fully tested in the protocol how all those pieces will work, where that information will come from, that kind of thing.

So that's something I would like to address in the pilot and take some time to figure that out, and to test all of that. I think an IANA registry of where to find the RDAP servers is probably not the right kind of system, so some kind of SRV record or something else is appropriate, maybe a standard location, I don't know; these are questions.

FRANCISCO ARIAS: This is Francisco Arias from ICANN org. Actually I think that question was already considered and I'm looking at Mark Blanchette, hopefully he can jump in here, but I believe it was already considered as part of the RDAP protocol, and it was decided not to follow, and instead of that we have the bootstrapping mechanics that allows for that. There is already a solution in the protocol, that's what I'm trying to say, for this.

EN

REG LEVY: This is Reg Levy from Tucows. Francisco, can you confirm whether or not participation in the pilot program requires access to live data?

FRANCISCO ARIAS: I'm sorry; I don't understand what you mean by live data?

REG LEVY: So we're thinking of building out a database that can test RDAP from our side, from the registrar's side, but we don't want to necessarily publish it until we're sure it works properly, so it would be using non-live data; it would be data from a couple of days or weeks ago.

FRANCISCO ARIAS: Oh, okay. I don't see an issue with that, I don't think there is any requirement in terms of the data being up to date or anything like that. One of the points in the proposal was to not have more requirements than comply with the standards. It's a pilot at the end of the day say, "We are testing things," so I don't think there is an issue with not using the most up to date version of the research and data.

REG LEVY: Great, thank you.

MARK ANDERSON: This is Mark Anderson. I agree with that, I'd just point out when you announce your pilot on the RDAP pilot page, there is a column for implementation, features and limitations that might be something you note in the limitations, "Hey, this might be stale data."

KAL FEHER: Kal Feher. One of the observations I'll make regarding people's comments is that there seems to be a lot of client-side concerns. So how's the RDAP pilot accommodating client software development?   It seems to be almost incidental and implementations that we have observed are generic or perhaps inappropriate for some of the solutions that we're looking for, especially when you're talking about tiered access to data.

                        And we're going to be providing any software writer basically a moving target as we change our implementations and test things. At the end of this pilot, do we hope to have good client-side software, and I think that would cover off some of Maxim's concerns regarding formatting and some of the other comments

that I heard about consuming the service, ultimately you need a client to do that.

MARK ANDERSON: No one else to take that one? I think that's a great point, and the only thing I can add to that at this point is one of the people participating in the pilot is a domain research LLC and they indicated that they're going to write a Golang client library for accessing RDAP services, so as Maxim pointed out, there is the readability of RDAP in its raw form isn't great, and so the development of client software certainly could help improve readability, but I think from a direct standpoint, that's maybe beyond the scope of what we're trying to accomplish here.

I think our focus is on the registry and registrar technical implementation of the services themselves but I would hope certainly that that would spur development of client software that would read those servers.

JIM GALVAN: Mark, may I?

MARK ANDERSON: Go ahead, Jim.

JIM GALVAN: So actually I want to just sort of turnaround; I want to add to that response in a slightly different way, and say that this is actually an intended feature. It's not a limitation of this pilot or anything else. I mean one of the main features of RDAP is the fact that you now have structured data, unlike Whois which is just dumping an ASCII blob at you, okay? You now actually have something that allows you to do interesting things and to build interesting clients.

Many browsers will actually parse that output directly, and at least give you some kind of presentation of it, but it's almost trivial to add a format that makes nice, pretty output, and it is intended that it's a feature of RDAP that various kinds of clients will come into existence to do interesting things with the data. And one of the more interesting thing you might have imagined is the ability to translate or transliterate some of that data, because you have now nice structured data so you actually have a way to work with it with those kinds of tools, if you are so inclined. Thank you.

MARK ANDERSON: Thank you, Jim, good point. Go ahead.

UNKNOWN SPEAKER: [Inaudible] coming back to the discussion of bootstrap, one thing that the speaker on the other side said about for the end users getting a central point, could be done by the registry but kind of a shadow registry just for the pilot, don't bother about IANA but just create the registry on your pilot page, and then people can start working on it.

The other caution was about SRV records; the bootstrap discussion within the ITF working group took many cycles and many meetings and there were no at the end final good solution; there's pros and cons in each and it depends.

So I don't know if there's a mailing list for this group, but I could send the actual URL with tall the pros and cons on each and what was good and bad for each side of one of the requirements was to actually make sure that we could run something on a tinny client that doesn't have access to any DNS, therefore cannot send SRV records for example because you are inside kind of a browser or some kind of restricted environment; that was one of the requirements.

Anyways, I'm not sure I want to go into the whole discussion, it was a long discussion but I could send to anybody a URL of the discussion. You may agree or not, but at the end that was the solution chosen by the working group.

MARK ANDERSON: Thank you. Francisco, did you want to get in the queue?

FRANCISCO ARIAS: Yes, thank you. I just wanted to mention there are a few clients, open-source RDAP clients available; at least four, I'll post the links into the chat. The RIR Registry for North America, and a leader in developing RDAP, they maintain a client and there is CNNIC, CentralNic and DNS Belgium that also offer RDAP clients readily available.

UNKNOWN SPEAKER: Thank you, and thank you also, Jim, for clarifying regarding the purpose of the RDAP trial, and that does allow for rich client development. My only concern really around this is that those clients need to support whatever authentication scheme ultimately settle on, understanding that of course we won't settle on that during the pilot, it would be great if there were enough examples, and we had some outreach to some of those developers, and I don't know who does that outreach and whether that should be in this pilot or not, so that they could at least have something useful as soon as we end up solving Whois, which is very soon.

JIM GALVAN: Okay, Jim Galvan again for the record, when you express concern about support for authentication schemes, I guess part of what I'm thinking in my mind is, this is fairly well oiled technology if you will. One of the things you get with RDAP is you inherit everything that you know about authentication in the web. It's all got http behind it, and you've got your web server kind of technology behind it and all the tools that come with that, so the credential management piece of this pilot is really just about understanding the presence of certificates and using them as your login mechanism.

I guess I don't want to overstate the perception that there is a significant problem here, because we are using technology that we should all well understand and know how to use; this is just about overlaying a particular framework for the certificates on top of that. So it's just a certificate-based scheme; there's not a lot of schemes to work with here. Does that make sense? I'm not trying to confuse anything, and I hope I'm helping, I'm trying to soften a little bit what I thought I perceived was a real concern about how this is going to work. Thank you.

UNKNOWN SPEAKER: Thanks, Jim. No, I wouldn't say I'm concerned, I'm just keenly aware of some of the limitations and trying to prompt perhaps this group to do some outreach to those software developers so they can get involved. So I take your point; it's not a client-side disaster by any means. It is a much bigger solution than Whois, so if we could just continue some of that reach out and if there was a way of providing feedback regarding client-side performance as part of this pilot, like perhaps some feedback to say, "Hey, perhaps this client is not working correctly for a particular implementation or perhaps a score card or something like that."

JIM GALVAN: So Mark, that actually brings up -- this is Jim Galvan again, something interesting that he just triggered me to think about. Our implementation as we continue through the pilot and we add features and the software changes over time, there's a version of it that'll be out there that people can always go use and look at, so an interesting point here is; is it possible?

General people could certainly go and look at it and use it and do queries to it, and I imagine it'll probably be able to query many of the other implementations; they'll be versions of them that are up there and considered in alpha or beta mode, whatever you want to call it. It would be interesting for the pilot

to find some way to give people an opportunity to submit feedback on what they're seeing on any given implementation.

I mean, maybe each of us who puts a server out there should also provide some way for users to comment. I mean, I don't know if we should collect it individually, or there ought to be some central place, some feedback can be collected, but I think that's something that we ought to think about here as we get organized and see what we're doing.

MARK ANDERSON: Thank you Jim. I have a couple of people that are in the queue so I'll get to you. I'll just, from a VeriSign perspective, we did put an e-mail link on your pilot page, just for that purpose; to solicit feedback, but I think your point is important. If people are just sending feedback to one implementation, there's not an opportunity for all of us to take advantage of that, so that's maybe something we can take back as a group and consider how to provide a better mechanism there. I think the gentleman over there whose name I don't remember.

GREGORY MOUNIER: It's not too difficult; Greg. Again, I could put practical questions for the end users and for those who would like to be able to be testing your different platforms. So first of all, ICANN is not

going to support a little bit more in terms of getting the feedback of the end users, so we have to send our feedback potentially to each of you, is that what you're saying?

MARK ANDERSON: I think what I'm saying right now is that we hadn't really considered that, but I think that's a great point, and maybe Francisco is looking to get in the queue on that one.

FRANCISCO ARIAS: Yes, thank you. This is Francisco from ICANN org. As a matter of fact in the pilot page it says where to send the feedback, so besides the specific implementation feedback mechanics that every single provider has, like VeriSign in this case, ICANN set up, or it's asking people to use the gTLD Tech mailing list, the details can be found in the first link of that page to provide feedback in general about the pilot, so that mechanics is already available.

GREGORY MOUNIER: Can I do a followup? And so another question, very practical again, in terms of Afilias so VeriSign, are you asking users who want to test the platform to login and to provide some credentials, and do you also log the search that we might be doing?

MARK ANDERSON: So our pilot at this point does not include credentialing or authentication, that's functionality that we would like to test out and try as part of the ongoing pilot effort. I don't have answers yet to those questions, I think those are things we'll have to work out as we go along, but I don't know if Jim or Stephanie wants to jump in on that one.

JIM GALVAN: I'll add and say the same thing. I mean, our intent here is to test credentialing and we're actually going to set up kind of a mini CA for ourselves, and we expect eventually to ask for people to want to get some certificates so that we can set up some interesting profiles for different people, just to test how things work. So yeah, answers to that are to come.

The current version that's out there does not support any kind of authentication on the access to it, but as that comes available we'll make visible that when we're looking for volunteers we'll look for those too, and I'm sure it'll all be visible in these forums, what different folks are doing.

MARK ANDERSON: Thank you. Does that answer your question?

GREGORY MOUNIER: And on the logging of the requests, something you haven't thought about? You know, we're law enforcement so we'd like to help of course, but of course we don't want our request to be logged.

MARK ANDERSON: Fair enough. We certainly for VeriSign's pilot we haven't considered that yet. So I don't have an answer to that, I don't know if Stephanie or Jim want to answer that?

STEPHANIE DUCHESNEAU: I would have to defer to Brian; if he's still on the bridge.

BEN MCLLWAIN: Hey, Ben McIlwain; engineer at Google. Very interesting question, I'm sure everybody would agree it's technically possible. The real answer is, would we want to do it? And I don't speak for the policy people, I just speak for myself; probably not.

JIM GALVAN:	Yeah, I guess I have to apologize; I don't understand the question that you're asking, so I'm hearing people say, "No," but I don't know what they're saying no to.
GREGORY MOUNIER:	Very practically; if I ask some of my analysts to test, and to go on the platform and do a few queries and stuff, if I want their buy-in, I have to be able to tell them, "No, Afilias is not going to log all of your requests for research purposes," for instance, or just to improve the platform or whatever. You see what I mean; we don't want to have our research and our tests -- I don't know, do you want to help on this one?
BEN MCLLWAIN:	Yeah, I think he's basically asking if we're going to save our search requests, and the things that we query, because we're querying people that we wouldn't want anyone else to see.
JIM GALVAN:	No, you know, it's just a test thing. Whois is still there, I mean if you're playing with this thing, we assume you're just playing with it for whatever purpose that you're going to be doing that for, and you're not using it as any kind of authoritative thing; we're not making any promises about the data.

I mean, it is going to turn out to have certain -- it's going to be a live reference for us, but it's not a replacement for Whois or anything else, and I certainly don't have any plans to keep anything around longer than I need to, to do the testing and developing that we're doing.

MARK ANDERSON: Is there another hand on this side of the room that I missed?

UNKNOWN SPEAKER: If I could just make a comment on that conversation; that sounds like that's a user requirement that probably needs to be considered for implementation in the future, so it's interesting that we're receiving these.

Is there a way we can capture this kind of feedback within policy or direct that to policy, because when it comes time to implementing RDAP, if that is a genuine concern and none of us have truly considered it, implementations, that might be an issue?

MARK ANDERSON: Joe, I'll get to you in a second. I'll say a couple things; first, the session is being recorded, so we do have a record of it. But also say, many of us participating in the pilot are also on the Next-

Gen RDS PDP and I think there's a hope that some of the learning's and lessons we get out of this pilot we can take back to the Next-Gen RDS PDP and apply there, so hopefully yes. Jim, go ahead.

JIM GALVAN: So thank you, Mark. Jim Galvan from Afilias. My gut reaction is there's nothing special going on here; it's just another service that we already provide. RDAP is just a replacement for Whois, no change to data retention or policies at all. I mean, you're certainly subject to whatever downstream policy changes come about, but this is not a special case and I don't expect to make any changes to what I'm doing unless I'm told by some other policy somewhere along the way. Thank you.

UNKNOWN SPEAKER: Yeah, I think it's fair. I think just something to consider, like if this came out down the line and then law enforcement were logging into something, there could be the assumption that things are being saved, and that's where some of that clarification would be needed.

MARK ANDERSON: Thank you. Go ahead.

OWEN DELONG: I think the difference here, number one; there's no reason for you to believe that your Whois queries aren't being logged necessarily, so if you're concerned about that you might want to think about that. But number two; I think that the difference that they may be want to be more concerned about with RDAP versus Whois is that they maybe be authenticating to RDAP and so the query log may not only identify the address the query came from and what was queried, but also who queried it.

And I can understand why various entities, including law enforcement might be concerned about that and so I do think it is something we probably should consider, but I think it's more of an issue for the RDS PDP because I think it's more a policy question on data retention, data privacy, data access than it is a technical implementation question of whether we include the ability in the server to log a query or not.

JIM GALVAN: Jim Galvan again, and I want to add on to what you said and make the following suggestion; it occurs to me that that kind of issue would be specifically addressed when the template response is being developed for that particular family of credentialed users.

One of the requirements might be that I'm going to have this credential, and by the way, if you're ever dealing with this particular credentialed user, you shouldn't keep any logs or whatever the rules are about the logs for that. So yes, I agree it's a policy thing and I also think that policy could be different based on what credentialed user is accessing it. And so I expect that question to be one of the check-boxes in the development of that profile and template response. Thanks.

MARK ANDERSON: Real good questions, thank you. Just sort of a quick time check; we've got fifteen minutes left. I think this has been a great conversation, very useful and productive so I'm inclined to let us keep continuing on this, but if there are no more hands… can you jump back to the slides?

FRANCISCO ARIAS: I'm having issues here. Antonietta, would you please help me? [AUDIO BREAK]

MARK ANDERSON: So sort of overcome by events here a little bit, but I mentioned earlier that we really want to encourage participation in this pilot. If registries, registrars and ICANN are developing a new

profile in a vacuum, we're dooming ourselves to failure from the get-go. So the more participation, the more feedback we get in this process, the more likely we are to iron out some of the issues and gotchas ahead of time. So that's certainly part of what we're doing here today, and so I greatly appreciate the feedback and input in the discussion we've had. But we do want to encourage that to continue throughout this process. Next slide please.

Again, overcome by events a little bit, but I'll just touch on these. At our first meeting I think it was two weeks ago, we made the decision that as a starting point we'd use the profile that staff developed almost two years ago for how to implement RDAP and RDAP is a protocol by itself; it doesn't solve or do anything. There needs to be some kind of additional direction on how to implement it. And the profile that staff developed is a starting point for that, and so we're going to use that as a baseline and go from there.

I added a bullet point on there for discussion point on one versus multiple profiles, we may or may not have time to really get into that today, but I'll go ahead and sort of tee it up. When we were developing the proposal to have an RDAP pilot program, it was suggested that we may end up in a situation where there are multiple profiles instead of a single profile.

And the example was given to me, was that we have brand registries that have perhaps a different use case for how they might want to implement RDAP, so there may be a need, it might not be a one-size fits all profile, so we at least wanted the flexibility to consider the possibility of having multiple profiles instead of one profile that every registry implements.

I'll pause there if anybody has any comments, thoughts, input they would like to give, we'd welcome it, but this is certainly not a decision we're looking to make here.

JIM GALVAN: Jim Galvan for the record; I think that multiple profiles -- let me start differently. When we say multiple profiles here, are you talking about multiple default profiles, or what's going on? And I guess the context for my question is that I think multiple profiles are a foregone conclusion.

I mean, we already know the answer to that question, and I base that on the fact that different families of credentials will have different profiles, it will create responses, the profiles represent the template for the response that you're going to get based on whichever credential family that you happen to be in. So I think are we talking here about just the default profile as being one or multiple, or is there something else going on here? Thanks.

MARK ANDERSON: I think all of the above; I think at this point we're still trying to work that out. I have no pre-conceived notions there on what we end up with. I tend to agree; I think it's likely we'll end up with multiple profiles, but I think that's all considerations we need to consider throughout this pilot.

UNKNOWN SPEAKER: Can we make this conversation a little bit less abstract, like what specific multiple profiles are we envisioning? Like are there some examples?

JIM GALVAN: So sure, I'll just make some stuff up. You know, if I'm a law enforcement person asking for a particular profile, I'll probably tell them everything I know about that domain name, so I'll give them the name, and the postal address and all the contact information that I have. It may also be that other kinds of families, it might be the case that I should give them only one piece of contact information; maybe they're only eligible to get the e-mail address for whatever reason, it might be someone.

So when that credentialed user asks about a domain name, I'm not going to tell them everything about the registrant; maybe I'll

just give them the name and an e-mail address. I won't give them their postal address, because that's regarded as being kept private. There might be another kind of credential, perhaps in a generic user case, not even a credential. I mean the default profile might be that only going to give you thin data and you're not going to see anything, so you're only going to see a registrar ID and that the domain name exists in the NS records.

Other scenarios might be; a particular credential might allow you to distinguish between a person versus a legal person, I'll just throw those terms out there without really defining them, and hopefully people understand that. And so if you understand that distinction in your database, and they're asking about a person, you may not tell them because that's a privacy violation, but if it's a legal person, you will tell them because some countries demand that kind of a distinction.

So that's what I mean by profiles; it really just determines what data that I give back and I just don't even know. The families will somehow as part of creating their family of credentials, they will define a number of interesting things. The profile of response that they get, there might be rules about logging that go along with it, things like that, data retention policies about what's happening. So that's what I mean by a profile; all of those parts. Does that help? Thank you.

FRANCISCO ARIAS: Thank you, this is Francisco Arias from ICANN org. So perhaps it's just the terminology thing, I guess the profile that we have published in the ICANN web site it's more than just the set of data that you see, which is what I understand you're talking about, and in that sense I think the term in the RDAP RFCs, you have the full response, and then you have redacted responses that vary depending on who is asking for the data or some other policy decision. But I think when -- at least in the terminology that we were using when we were talking about profile, we were referring to something more than the data. It was about what functionality will be available, and certain elements that need to be there.

For example, the remarks elements in the RDAP response, what they need to have and other technical elements that are part of the RDAP profile so that you have a common response, or format of the response in RDAP within the gTLD space, at least that's how we use the term. I'm not suggesting that's what needs to be done here, of course a decision can be done to use it otherwise, but just pointing out what we use. In some way we were using the term profile to refer to the subset of the functionality that was to be used, and when talking about the responses, we were

talking about redacted responses to signify the different sets of data points that a given user would have access to.

MARK ANDERSON: Thank you. I believe there's a comment down there?

MARC BLANCHET: Yeah, I was going to say similar to what Francisco said, which is most of the profile is making sure that when you send some data, this is what the others should expect and https transport and all kinds of stuff, it doesn't really say that much about if you want to remove some data from in your response, so the current profile may have some changes, small changes to accommodate, what Jim said, but roughly speaking it's not about what did you send, but essentially how did you send it, how you format it and making sure that for example, IDNs are properly sent and things like that, all those kinds of stuff.

MARK ANDERSON: Thank you, any other comments? Alright, looking at time, we're at the five minute mark here so I'll talk about a couple of logistics items. Currently as I mentioned, we're starting with the profile that ICANN staff developed, it's in a Google doc format, I was going to ask that we get into a sort of logistics of how we

start drafting, but in the interest of time maybe that's best left to a follow-up meeting. Antonietta, if you can go to the next slide?

I want to touch on this, it's certainly the hot topic here this week, I assume GDPR needs no introduction. Everybody knows what GDPR is, great. I think we're not really going to solve GDPR here in this group obviously in five minutes, but I'm sure you've been in meetings where there's been discussions about getting to GDPR solutions and I think there's a lot of hope among the community that people are going to start presenting IDS and solutions for GDPR and that's an opportunity for this group; those of us running pilots to actually test out some of these ideas.

So as the registries are putting together the proposals for ICANN staff to run an RDAP pilot, we were certainly keenly aware of the GDPR elephant in the room so to speak, and we knew that there would be the possibility at least of doing some work in parallel as proposals for GDPR come out, we could potentially test some of those, and again kick the tires, see what's going to work and not work.

I'm not aware of any magic silver bullets for solving GDPR yet, but that's something as a group we'll certainly be keeping an eye out for, and hopefully we'll have an opportunity to try to some of these things out before May 2017 gets here. So I'll just

sort of put that flyer out there so everybody is aware. Next slide please.

So I mentioned at the top, this is a new group, this is only our second meeting ever so we're getting our feet wet here. I made the mistake at a previous meeting of trying to figure out a schedule for regular meetings in the meeting and that didn't go well, so what I'll say is we'll send out some e-mail and try and figure out logistics for a day and time that'll work well for everybody for scheduling regular meetings, but if we're going to keep momentum going, keep people talking and communicating, we'll need to have some sort of regular meeting scheduled. So keep an eye out for your e-mail, and we'll see about getting that scheduled in the very near future. Next slide please.

And that's it. You have two minutes; any thoughts, additional comments? Stephanie?

STEPHANIE DUCHESNEAU: Just wanted to say thank you to ICANN and to Francisco specifically; I think when we all went around the room in Copenhagen, asked where we wanted this issue to go, a lot of us asked for exactly this; a flexible pilot that allowed us to sort of test it out and have the kind of conversations that we're just

kicking off today, so I think it's a great example of collaboration and I'm excited to kick off work with the rest of the folks in the room.

JIM GALVAN: This is Jim, I just want to second that, too. Very important point; we worked hard to get to this place, and appreciate ICANN working with us to get us here and looking forward to some real progress and substance, thanks.

MARK ANDERSON: I was going to second that, but I'll third it. Thank you, Francisco; your support on this has been excellent. Any other last shouts from anybody? Alright, thank you very much, I appreciate everybody who attended and all the comments and participation. With that, you can end the recording.

[END OF TRANSCRIPTION]

ICANN 60
ANNUAL GENERAL
ABU DHABI
28 October–3 November 2017