ICANN
Transcription ICANN63 Barcelona
GNSO – NCSG – Can formal standards simplify 3rd party access to registration data?
Session 2
Sunday, 21 October 2018 at 15:15 CEST

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

Stephanie Perrin: Well we'll have to get out the cattle prods and round people up again I guess. But I think we should get started. Over to you, Ayden.

Ayden Férdeline: Thanks, Stephanie. And just - has the recording started? Perfect. Thanks very much. Welcome back. It's good that we still have some people in the room and hopefully a few more people will enter shortly. But next on the agenda we have Mark, who's going to be walking us through Microsoft use of privacy standards. Thanks.

Mark Svancarek: Shall I? Hello, everyone. I'm Mark from Microsoft. You can say my name any way you like. Usually I go as Mark S-V. So I don't have any slides today but I'm going to talk about what Microsoft is doing in privacy and security and to that end I'm putting a link into the chat room and this is the Microsoft Trust Center specifically the privacy hub portion of it.

The Microsoft Trust Center is a resource that we have for our commercial customers; retail customers and end users are directed away from that to more specific sites that are more appropriate to them like, you know, what data does the Windows Operating System collect and what control do I have over it, that's more of a consumer question even though it applies to corporations. Corporations are usually concerned about bigger issues such as, you know, how do - what is FedRAMP and what are the privacy protections on my data.

So in general, our - there's sort of two categories of privacy controls that we have. Some of them are based on certifications and standards and others are based on best practices. So in the Trust Center you can see things like here is the support for ISO 27000-18 and you can get more details on it, things like you know, is this based on the attestation or audit? It's based on audit. How often do we get audited? Things like that.

And you can drill further and further into that to learn things like what the list of processors that we use, how do we use your data? What are Microsoft's responsibilities? What are your responsibilities? And in some cases there are standards where we conform to the processing requirements for the standard but really the responsibility of the data is the customer's responsibility for the data.

One of the things about standards is that everybody's got one. And so we try to work with customers all over the world and governments all over the world. In recent years governments have been proposing a lot of security standards and privacy standards and we try to get in front of that and show them, you know, we support the following things, you know, here are the standards and certificates that we support. What you're proposing is 90% the same as ISO or is significantly like this other thing and try to guide them towards some sort of standardization because we think that's generally better for everyone, usually these additional - these additional standards don't add a lot from our opinion in terms of control or actual protection but they complicate things for everybody, right?

So if you need to certify within your country, now you have to come up with your own set of certifications and, you know, maybe it's the Ministry of Information who does it or maybe it's a third party but those people can't really do that anywhere else so it’s really inefficient for everyone involved and so we try to you know, guide people towards a common set of standards.

And so on this site, let me go back to it, sorry. So we have a bunch of explanations about, you know, our practices, how we adopt privacy standards, how we build privacy into our services, where we have contractual commitments to back our practices so like 27-18 is contractually backed for instance.

And then there’s more details about 27-18, EU model clauses, the My Number Act in Japan, the Argentina Personal Data, the Canadian Privacy Data and there are others but these are the top ones that we have on our site. There’s actually a pretty big heat map that we use internally to make sure that we are world class in all of these certification areas, you know, according to our own mission, but also relative to our competitors.

So I don't know if you are aware of how controls are put into a cloud service. There are different layers of controls; there are physical controls, such as all the servers are in locked cages and only certain people have access to them. There are process controls such as in a China sovereign cloud all the IT people have to be Chinese citizens. So some of these controls are standardized and some of them are specific and it’s particularly in places where there’s data sovereignty laws. Those facilities are completely separate and have additional sets of controls on them.

Let’s see, so again on the Trust Center, I suppose I should have created some slides, sorry. We talk about how we handle government requests, we show our position on the Cloud Act, what are our own privacy standards such as we don't allow direct access to customer data, we redirect all the law enforcement requests to the customer, we don't give access to platform encryption keys, here’s our position on intellectual property and stuff like that.

And then also we offer services to our customers. As I said, there are a lot of the responsibilities that are actually the customer’s responsibilities. You know, so we can provide a framework for you but if you leave your service unlocked and people can get in or if you haven't partitioned your data properly so that one set of people can have access to things that they don't need, you know, that's kind of on you. But we do provide a lot of resources e-books, training, we have technical account managers who will come on your site and train you and we help you with audits and things like that.

We don't expect that all customers come in at any sort of level of expertise and so we try to make it - make it easier for the customer to be safe. And if the customer requires their own certifications - and this is true for instance if PCI so the payment card industry has a certification and most of those responsibilities are not on us as a cloud provider so we achieved our own level of certification but really it’s about making sure that any customers who are running their business on our cloud, you know, these financial institutions, that they know how to be safe on our cloud because it, you know, reflects badly on us all if you say I’ve adopted Azure for my business and I got hacked.

So I think that’s the general introduction to what we’re doing. Regarding individual privacy, a lot of it comes down to telemetry. And you may recall a few years ago that there was a controversy about Windows 10, you know, collecting a lot of data. And that was because we hadn't been very clear - we had convinced ourselves that we were, some of us said that we were not and of course it took a crisis to get everybody to wake up about it.

We were not clear about what data was being collected by Windows, how it was being used and what you could do about it if you elected not to do it. And so, you know, we had to break it into chunks like here’s the data that we collect just as part of running the operating system to make it better like if there are - if you have crashes we collect the data on, you know, which apps crashed, which processes crashed, how often they crashed, if it had something to do with a particular piece of hardware, you know, maybe that points to a kernel mode device driver or something like that.

Then there was other data that we collected that was actually personal data, so I would like, you know, settings, like I would like my settings to migrate across multiple machines. So at some level that’s personal data, maybe not identifiable but, you know, it’s related directly to you.

And then there is personal data like my Microsoft account, so at some point I signed up for a Microsoft account which allows me to access services, that's personal data. And so we had to be very clear about here are the various categories of data, here’s what’s collected, here’s how you turn it off, but here’s what the benefits were that you're going to lose if you turn it off so that customers can make a better choice. And of course there were a number of customers who turned things off and there were some businesses who turned the telemetry off too. And, you know, they get different levels of service, but that was their choice.

So I’m wondering, are there any specific questions that I can address to make this more interactive? I mean, I could keep talking about, you know, what we do. If not I could talk about, say, DPIAs or something like that?

Stephanie Perrin: That would be interesting.

Mark Svancarek: Okay, so I’m not actually allowed to show you one of our digital privacy impact assessments, I’m not actually even supposed to show you the table of contents; we reserve that for auditors. But I can…

((Crosstalk))

Mark Svancarek: Yes, I can speak in very general terms about it so I’m going to open it up on my desktop and just talk to it. So what I’m looking at right now is the DPIA that we generated for our internal usage of Whois data. And this is a 20-page document. And it’s broken into chunks; there’s a chunk for privacy managers, there’s a chunk for frontline attorneys, one for the data protection officer, and one for corporate, external and legal affairs.

And so it goes into the list of all the business point of contacts, so there’s a digital crimes contact, a threat intelligence contact, anti-piracy, trademark infringement, SSL certificate administrator, and corporate domains administrator. And then we go through with the description of all the processing that all those people do, where the data is stored, how it’s controlled and like that.

And so we've done many, many, many of these throughout Microsoft. Our adoption of GDPR took us about two years. We had to go through, you know, as you can imagine, everything, so all the human resources records, we had years and years of records of, say, customer escalations, you know, so corporate customers have a different escalation process than retail customers, you know, so if you - if you're a retail customer and you call up customer support, that’s different from I’m X company and I’m having the following problem.

So those are logged in different systems and so all these various systems had to be scrubbed, you know, for privacy. And some of them couldn’t be salvaged and we wound up throwing out a lot of data, baselines and historical data and things like that just because the system itself could not be made compliant. So it was kind of tragic and painful but these are the sort of decisions that we had to make.

And at this point now we feel like we are extremely compliant both internally, you know, for our employees but also in regard to how we handle our customers. And that’s a separate issue from how our cloud services are managed; we feel like we're extremely compliant there as well and in fact my vice president has blogged about the significance of GDPR and how we feel it’s a big step forward for the world and we’re hoping that this can be a differentiator for us in fact to say, you know, we can be the most privacy compliant cloud provider in the world.

Stephanie Perrin: Can you just tell us which template you used maybe?

Mark Svancarek: Actually I specifically can’t.

((Crosstalk))

Mark Svancarek: Yes, I know, so some of these questions are going to be kind of boring because, no. But I’ll try.

Stephanie Perrin: Which vice president blogged that?

Mark Svancarek: Julie Brill.

Collin Kurre: This is Collin Kurre for the record. I was just wondering - I found it really interesting that you have listed these different data protection laws, so in becoming GDPR compliant, were you able to kind of revolutionize the way that you handle user data in a way that you think will be compliant for other data protection regimes? And I say this obviously in the interest of perhaps informing ongoing ICANN policy development processes to assist them in being more future proof should conflicting or, you know, should things - other legislation that doesn’t - isn't a one for one match be kind of leveraged on this…

((Crosstalk))

Mark Svancarek: Yes, there’s an internal initiative that’s called Next Generation Privacy, I’m not sure how much of that we've made public but I could get you more information about it. And that is the effort to be future-proofed, so we have a number of these initiatives. There's a crypto-agility initiative that's similar to that, you know, because security is not the same as privacy but they lean on each other and so, you know, we try to figure out how do we become more agile and future-proof both technically and within our internal processes?

And so the initiative that's related to that is called Next Generation Privacy, which is bigger than GDPR and will be ongoing forever. So 15 years ago or something like that we created the Trustworthy Computing Initiative, that's where we first started getting interested in privacy; now we're into our second phase.

Man: (Unintelligible), for the record. So from your experience in implementing privacy safeguards standards in Microsoft, what do you think would be the biggest selling for GNSO interest in implementing formal standard for simplifying the third party access to registration data?

Mark Svancarek: Well, in all cases I would suggest that you do, you know, the standard things like make a data inventory, make a data flow diagram, figure out how that's being managed, document it all, in a DPIA, have someone who is an expert on this topic review all of those things for you, you know, don't try to learn it on the fly by yourself would be my advice. Regarding Whois specifically, you know, right now we're in this situation where every request is subject to a balancing test at the receiving side of the request. And I can't really give you any specific advice there. Right now it's a heavy process and, you know, you just have to do your best and manage your risk appropriately to your situation.

But in terms of getting yourself into a good state generally, you know, separate from Whois, those steps are pretty standard. So really make an inventory, what is all the customer data that you have? And you might be surprised, that's why you have to do the inventory because you've probably made assumptions about what you have and where it is and who's using it and who has access to it. And until you really start to do the process you won't know if you're compliant with anything or not or if you have the ability to secure it or not.

And in our case we had to, you know, deprecate whole processes and systems in order to make ourselves internally compliant or if not internally compliant with something, to a level that we felt was appropriate. So there are many cases where we have our own internal compliance standards that are not based on industry standards but just values and targets that we have set out for ourselves.

Ayden Férdeline: We'll go to Stephanie next and then we have a question from a remote participant.

Stephanie Perrin: Thanks. Stephanie Perrin for the record. You've kind of half answered my question already. You are on the EPDP and so am I and you've heard me and Collin shrieking for privacy impact assessments since the get-go and even from the beginning before we - when we were doing the charter. Do you think there's any merit in trying to do one now and how complicated would it be if we did?

Mark Svancarek: Well, I mean, it seems like there are some things that probably have been done so when we talked to (Jaan), I got the impression, although he didn't say it explicitly, that at least a data inventory had been performed.

Stephanie Perrin: But it'd be so nice to get our hands on that wouldn't it?

Mark Svancarek: You know, yes, I mean, they're a transparent organization ideally and they have transparency initiatives elsewhere, so it would be great if, you know, if they could share more of the stuff with us.

Stephanie Perrin: Well…

Mark Svancarek: I am being candid though that there's a lot of this stuff here that I'm not allowed to share with you that I would have chosen to share but I was asked not to share and so it might be inappropriate for me to pass judgment on

ICANN's corporate decisions. But just speaking as me, yes, I would love to see that stuff. I would love to learn more about their stewardship of various sets of data. And I would love to understand where they need to build up their muscle in order to, you know, comply with any sort of laws related to data stewardship.

Stephanie Perrin: Well the fact is Microsoft, Stephanie Perrin again for the record, is a corporation, your privacy features are a competitive edge. I can understand why you wouldn't want to release anything. However, for those of us toiling away as volunteers on the EPDP, it does sort of beg for a document request and we are planning on doing one because if there's an inventory - if there's a map, why not show us, you know? Thank you.

Mark Svancarek: Yes, and at some point, you know, if there ever is any sort of an access model, if ICANN is touching any part of the data in an access model, I mean, they're already touching it, you know, for compliance purposes. They're going to have to perform these activities anyway. I personally would like to see them done earlier than later whether or not the contents of these assessments are shared with us in the community or whether they're only shared with, you know, third party certification organizations, you know, that can be a discussion that we would have.

As I've said, I would love to see the stuff myself, but if it's decided that they have to be managed more tightly and reviewed by third parties only, I would just, you know, would just like to see, okay on this date we submitted this to this party and here is the report that came out, I think that would be a good step forward.

Ayden Férdeline: We have a question from Steve DelBianco. I'll read it verbatim. "Mark, you mentioned your IRM, who determines which user gets access rights to data in IRM? Is it similar to federated access?"

Mark Svancarek: Okay. So in the case of IRM, this is a very user-friendly tool. So in the most trivial case, I create a document and then up in the ribbon at the top I can grant specific permissions, I can grant permissions - the most typical use is this can be read by anyone at Microsoft but it can't be shared outside. That's the most typical case and in that case the author would set the controls and then there'd be a very broad readership. There's other sorts of things like you can read it but not write it, you can write it and edit it, you can send it to someone else, you can limit it to certain security groups.

In other cases, the document might go into what we call a high business impact SharePoint site, HPI, and in those cases, these things are locked down at some confidentiality level and there the confidentiality level is being determined by some sort of an administrator within a business unit. They're deciding who is in the security group and what rights are allowed. And that would be something that where the author might not actually be setting - the author is probably not the administrator and so it's at a group level that the right are being decided. And in those cases it can be very, very locked down, maybe only a small number of people can see it or a single business unit or something like that.

And so relative to a universal acceptance thing, I think we have a lot of flexibility in how we impact this, I mean, implement this. Let me see. Could you clarify your question a little more, Steve? I want to make sure I'm answering the right question.

Ayden Férdeline: Just while we hear back from Steve we might take the next question from Andrew. Andrew.

Andrew Clement: Andrew Clement for the record. It's great to hear that Microsoft is really stepping up to GDPR and wants to use that as a selling point. Maybe slightly tangential to the discussion, but I'm interested in the question of where Microsoft stores this data and particularly European and whether that stays within Europe or whether it goes back to North America. I mean, obviously this is an issue in light of surveillance questions.

Mark Svancarek: Yes, so there are a couple different ways that we implement clouds. One of them is what we call a general multi-tenant cloud. And so many, many customers have their data within a collection of data centers that are distributed regionally or globally and the data is partitioned so that one customer can't see another customer’s data but they don't really care where the data resides. And so it gets moved around in some flexible and elastic way in order to maximize performance and reduce cost.

In other cases, we have government clouds and government clouds are much more locked down; they're usually within a certain region and usually there’s only a very limited number of tenants. And these tenants, for example, you might have the US Department of Agriculture in the same data center as the US Department of Justice. So it’s all the US Federal Government but there are multiple tenants within that and they would be segregated from each other as if they were separate companies.

And then there are sovereign clouds, and, you know, we discourage people from doing those because they're really, really expensive and inefficient so we have one in Germany, we have something with a different name in China, and there’s a couple other ones I think. And in those cases it’s very, very restricted. So it’s not just where the data lives, you know, at rest but where it can go in transit and you have to set expectations really clearly.

For instance, if you're a multinational corporation what are the limits on data sovereignty if you have a user who is trying to access data from outside of the national boundary? So you have to set the expectations up front. In some regions the authentication service might not be collocated and so you have to have the conversation and say, okay, well when you log in credentials are going to be going outside of the national boundary, are you okay with that or not because, you know, there are many implications on that.

And some regions can't justify having their own authentication center; it just costs too much. And so you have to set the expectation with the customer that, you know, some features may not be available in all places for practical reasons. But we do have the ability to lock down specific data to specific regions in many cases.

Ayden Férdeline: We received a clarification from Steve DelBianco. He asks if you could draw a parallel between IRM and RDAP.

Mark Svancarek: Okay I think what Steve is asking is, you know, I had made an analogy earlier about IRM and access to registration data. And so the way I was thinking of it is registration data has been collected somewhere and now it’s living at a registry or a registrar. And then someone needs access to it so say it’s, you know, different examples, cyber researchers versus Interpol. Someone has to accredit these entities at some level and assign them their access permissions. Depending on how this is implemented, this could be very coarse or it could be very fine grained.

And that assignment is happening at the accreditor. It can be processed at the contracted party but it’s not be assigned by them, it’s being assigned by the accreditor in their local region. This is just one implementation, there's a lot of ways we could do this.

But, you know, it could be done by the accreditor in the region and they can decide which rights are available in that region for instance, so this is similar to I’ve created a document and now it’s going to live on this SharePoint site and Joe the business unit manager has decided, you know, this is who can access the thing and what they can do with it. Can they - is it read/write? Can they email it? Can they do this? Can they do that? Or can they just read it?

And so you know, you could get into very, very complicated implementations, you could get into very simple implementations and those implementations can vary from place to place. And so I don't want to over promise any particular implementation because we're not there yet, you know, you could promise the moon, any of these things can be made, that's how software works but in actual fact some software is really, really hard to create and maintain and so it's always good to have reasonable expectations or maybe create things in an iterative manner.

Another thing about having different behavior in different regions is the same thing as the standards proliferation conversation that I had earlier. If everyone can agree on a couple of rules like open ID versus OAuth versus a couple things like have one or two standards, for instance, and then everyone agree on certain ways of creating rights certificates, you know, is that an XML that's attached to an X.509 certificate or something like that.

You know, if we could decide on just a couple of these implementations then people can start figuring out what's, you know, how to map rights and permissions onto those things without having to create really, really complicated interpretation systems at the side that has to receive the query and respond to it. But again, we're not there yet so, you know, these are just general considerations. We'll get there eventually.

Ayden Férdeline: Thanks, Mark. We're nearing the end of our time for this slot but we have time for one more question, Farzaneh, did you have something?

Farzaneh Badii: Farzaneh Badii, NCSG. Oh, so what I did have a question but then I'm seeing that at five we are going to discuss the law enforcement access to data from civil society's perspective, which I think what we are missing when we are discussing access and always we see that governments request to have access to Whois and domain name registrant data, well at this point we are treating the government as GAC, but later on which government is going to have access to personal information of domain name registrant based on the human rights - the internal domestic human rights violations because law enforcement accountability and checks and balances in many countries is nonexistent or they end up use of power can lead to serious danger.

And the fact that Whois was open for so long it - we don't know but it might have damaged a lot of things. So yes, but it seems like Benedict wants to speak so I just pass it.

Benedict Addis: I was just going to respond to - I will note that I haven't been a cop for four years. Law enforcement is very robust in talking about this issue and we - and there are some really good standards at Interpol, for example. Interpol, for those of you don't know, isn't actually a police organization, they don't investigate anything. They're kind of like a switchboard and they connect police forces in different countries together so they're really all about data sharing inter-jurisdictional so they fight with the stuff a lot for the last 100 years.

And what they say is you cannot use the Interpol network for certain categories of things, so religious, crimes against religion, crimes against speech, so on and so forth; they just say that's not in - allowed. And I think that we could draw on those rules and those standards in using in defining whatever access system - it’s not a free for all for law enforcement, and I think law enforcement is really comfortable with that idea. You can say that stuff, it’s fine.

Farzaneh Badii: Yes, just to follow up, yes, we have had conversations about this and it seems they understand of course, but the problem is that well maybe this is not the place to discuss this because it’s Stephanie's workshop so and this is like something that I have an issue I have with GAC, so I’m just going to stop here and…

Ayden Férdeline: I think we should stick to the agenda for now. So do we have any final questions for Mark? If not, thank you very much, Mark.

Mark Svancarek: Well thank you for having me.

Ayden Férdeline: And so, okay.

((Crosstalk))

Ayden Férdeline: Okay. So we’re just waiting for either Patrik or Greg Aaron or if there is someone else in the room who’s going to speak on their behalf.

Stephanie Perrin: Maybe I can - Stephanie Perrin for the record. Maybe I can take this opportunity to ask Rick some more questions about the RDAP implementation. The - one of the problems that are dealing with is there’s a - I think a misconception of tiered access and you hear the expression “layered access and tiered access.” And I understand that even though these aren't agreed definitions, tiered means different tiers that may form layers but they're also - it embraces the notion that you only get small data sets or limited data sets as opposed to access to the entire layer.

Whereas layered access is being interpreted as meaning there’s a layer here and once you proved that you are a - pick one - a cop, an intellectual property lawyer, a business interested in consumer protection, then you get the whole layer. Obviously we don't think that would pass under the GDPR because it’s not proportionate and it’s not limited and, you know, there’s a lot of extraneous data there.

However, figuring out how to do the RDAP queries while expressing the purpose and the limited nature, that’s my question, what can RDAP actually do about that?

Rick Wilhelm: Sure. So Rick Wilhelm for the record. In the RDAP Working Group, there currently aren't operating definitions of either tiered or layered. So therefore I can safely say that we don't have incorrect definitions of tiered or layered.

((Crosstalk))

Rick Wilhelm: Yes, exactly. So and to that end it's - and one of the reasons that we haven't even attempted to define those terms is that they are heavily depended on policy for the reasons that you even vague off the cuff, were sort of pointing out, they're difficult to sort of nail down, right. So we haven't, in the RDAP Working Group even attempted to capture those sorts of things, right. So there are - so in the chat Steve DelBianco types "The RDAP profile could include purpose of query for any accredited tiered access requests." So and then he says, "That would be logged by the RDAP server." That's a - so it certainly could, whether or not it would is a different - is also a matter of policy as Alex points out correctly.

There is a mechanism that's present for purpose in the RDAP protocol, I believe it's a short string capability and but what those purposes are have not yet been defined, it's just the - a bit of payload that's been defined. So we - so that also has to be worked out by some policy body as to what those purposes might be. So there's a bit of mechanism there where that could be captured after those purposes are defined, right. And then so there's a way for a query purpose to be communicated, right, and then acted upon appropriately. But the RDAP Working Group right now hasn't attempted to say what those might be.

After they are defined, presumably they would get captured somewhere in an IANA registry or something like…

((Crosstalk))

Rick Wilhelm: …something like that, now batting Manny Mota. So but that's an obscure reference. So although the Dodgers did win last night. The…

((Crosstalk))

Rick Wilhelm: So but right now the RDAP Working Group is just steering clear of all that stuff because someone else has to make those - we've got the mechanism in there for capturing the purpose and then decisions about what that purpose is is - we can act upon those - the RDAP servers can act upon those.

Stephanie Perrin: I mean, Stephanie Perrin again. It strikes me this could be immensely complex but as I mentioned earlier, I’m not geek, maybe it’s a lot easier than I think. But let me give you a couple of hypotheticals, if I am the, I don't know, the Canadian agency that looks after protecting endangered species and I’m looking for websites selling bear parts and turtle carcasses and things like that, would I have to specify okay, I found all these websites selling bear carcasses and I can give you a list of 400 of them, and then I would testify as to my identity as a regulatory officer in a given agency.

This isn't criminal, you see, this is not - it’s not easy. And every separate kind of investigation would have that complexity. And they may be looking for other things as well that would require a broader search. So you know, it’s a bit like trying to imagine a research search online before you’ve thought out what you're looking for, you know what I mean?

Rick Wilhelm: Yes, that’s a question for the - in and around the EPDP or its successor…

Stephanie Perrin: Ha.

Rick Wilhelm: …is sort of how that works. Really it’s the - and I’m successful successor. But that really - as far as how that request process goes and how that’s decided the RDAP protocol will implement that. There’s enough there or if it’s - if there’s not, if the policy comes down and there’s not enough there, then we will busily get to work extending it so that we can support it.

Stephanie Perrin: I have to say when I was on the EWG every time I asked Scott Hollenbeck these questions, and I’m not sure whether he was just trying to shut me up or not, but he would say oh, yes, we can build that; oh yes, we can build that. So

I'm really dreaming big about RDAP so let's hope it can handle this because it's a nontrivial problem, getting this away from wholesale access to specific queries.

Ayden Férdeline: We have a short queue forming. And please go ahead.

(Joyce Ling): (Joyce Ling), I'm a registrar (unintelligible) dotCom. Maybe, you know, I keep hearing about the pure (unintelligible) the layer (SS) and maybe we are making things a little bit more complicated than we should because look at the current Whois - public Whois, you already have the name server, the name creation date and some of them have the city of the registrant and the country, right? There's no email or no personal names. So if you're somebody want a tier, how many more tiers can you go? Get a personal name so that you are allowed to have only the personal name, so how could you get a personal name without the email and address?

So maybe tier and layers just give me the impression that oh my gosh, this is so complicated there, maybe things is much simpler than we thought it was.

Farzaneh Badii: I have one comment. Farzaneh Badii speaking. So what worries me about RDAP is that people keep asking, so can we do this with RDAP and that with RDAP and like all these like various design ideas that - but we should also know that RDAP might be able to do a lot of good things but it can also be developed to violate privacy or, you know, it's not - I think we should be a little bit more careful in how it's being developed. And a lot of your decisions are technical, I agree, but from time to time you might make some decisions that are based on policy or that you have.

Rick Wilhelm: So very fair point. The RDAP Working Group, as the group is going about it, is being as certainly as careful as we can to whenever we come across something that's policy to kick - defer it out to policy. The profile - RDAP profile documents were restructured earlier this year to be subdivided into implementation oriented thing - an implementation-oriented document and a document that is oriented towards absorbing policy changes so that that would be the area where policy-dependent parts would be explicitly called out and an area which are more technically oriented parts.

So we've worked to make it more apparent about the areas where policy would be influencing the document to where - and there's another section that would be - that we expect will be less impacted by policy changes. So we've done that in an effort to tease those areas apart from one another. That is an ongoing process. You're exactly right in that the profile - the RDAP technology itself could be used to throw access open - throw the (unintelligible) open and let everybody in and all the data out. If - that is a policy decision and it is a - it is just an access profile; or it can be an in-access profile. So that's entirely possible.

Ayden Férdeline: Alex followed by Marc.

Alex Deacon: Thanks. Alex Deacon for the record. Where should I start? So I think like all Internet technologies, you know, the technologies themselves are layered so RDAP kind of defines the data that goes back and forth. You need another layer on top of that that describes the authentication and the authorization mechanisms that will be used. And the one that was referenced earlier, OpenID Connect I think is a good one, it's one of many. For those of you who have read the BC and IPC draft accreditation and access model, buried in the back of that document - I'm sure it's dog-eared and you have it quickly accessible, is Annex I. It's a draft RDAP OpenID Connect profile that I wrote as an example of how OpenID can be used to convey claims around what the purpose of the query is, who's querying it.

We could, based on policy, we could add other claims that may be required for the recipient of those claims to process properly the request. And these requests are not, you know, blanket requests, they're domain name per domain name. These claims can be - will be conveyed if this profile or a version of it is adopted on a case by case basis.

And then to address the question by Stephanie, based on the authentication like the who, and the authorization like these claims, the what, the response profile based on a policy that's still to be defined can be actually quite granular. Right? We could define what gets returned based on the who and the what. And so I'd suggest you take a look at that.

I'm not - unfortunately I'm not and can't be part of the RDAP Working Group. I think that's limited to contracted parties, although I'll request again that that be opened up a bit. But I expect at some point in the future this group will be looking at how to - the technology that underlines the policy that we think will be coming down the road. And I agree with Scott Hollenbeck that these technologies I think can be quite easily used to implement an effective and secure tiered access or layered access mechanism.

Stephanie Perrin: If I could just do a lightening response now? Here, Alex, is a perfect opportunity for the IPC and the NCSG to join arms and work together and get on that RDAP Working Group. How's that for a proposal?

Alex Deacon: Maybe you've had better luck than I have in joining that group.

Mark Svancarek: So, you know, the reason that Scott answered all those questions with, "Yes, you can do it. Yes, you can do it," is because RDAP is just the protocol and, you know, so it's how you - the signaling and the transport and stuff like that. And it is a really well defined protocol so when I reviewed it I was very excited that I really couldn't come up with any use cases that couldn't be implemented with it. The use cases are not just implemented with the protocol as Alex said, there's other aspects to it, but it's a full featured well-spec'd protocol.

But what you want to pay attention to, and this is what Alex was saying, is the profiles. So the implementation profile is of all these features that can be supported using RDAP, here are the ones that we're actually going to build a server to use and some of us fight about whether, you know, X feature should be in this version or not. But the implementation profile defines which features are going into this version right now and then the response profile is kind of how are the features actually implemented in themselves?

So if you receive such a command how should you respond? And so you know, I review them with a couple of hats on, right. So UASG hat, how are IDNs returned, can you support a query that is a mixture of U labels and A labels in the same query for instance? And then also these other things like what kind of searches can be performed and stuff like that. And as Rich was saying, you know, right now, the kind of searches that can be done are exact domain name, you know, you put in a single domain name, it returns data about that one exact domain name. So if you want to look up 400 things, that's 400 queries.

Now you could - RDAP, allows you to do all sorts of searches. You could do reverse searches, entity searches, things like that. The protocol allows you to do that. There's - if you look at the implementation spec and the response spec, you'll see that those are not currently being considered and that's just part of a big list of things that are policy-related. So if a policy were to be created to allow or require such things, the RDAP protocol, which is very comprehensive and well defined could in fact support it, but they're not in the profiles right now.

So what you want to do is you want to focus on those profiles; you don't want to read the RDAP RFCs, you want to focus on those profiles and give feedback on them and they are somewhat technical so, you know, ask for help.

Stephanie Perrin: Thanks. We'll do that.

Rick Wilhelm: One last thing because I think we've got to move on, but I would be remiss if I didn't mention to help broaden people's standards that - understanding that

RDAP is not just used in domain name registries, there's a - it was jointly developed with the numbering folks and the numbering the RIRs, the IP address registries were actually hip to hip and in some ways ahead of the domain name registries particularly the folks that Aaron, Mark Kosters and such were doing this a lot and so the RDAP as one of the reasons RDAP as a protocol has such amazing flexibility, as Mark Sv was saying, was that it was built to accommodate not only domain name registries but also number address registries, so.

Ayden Férdeline: Thanks for that. So we are a few minutes over time at the moment so we might move onto the next item on the agenda. So Greg Aaron is now in the room, so welcome. Thanks for joining us. And if you would like to speak to us about what cyber security researchers would like, that would be very useful.

Greg Aaron: Thank you. I'm Greg Aaron. I'm here in my capacity today as a Senior Research Fellow at the Anti Phishing Working Group. And thank you very much for your invitation. It's good to be with you today. So I can tell you a little bit about APWG and its membership and what it does and APWG has been following what's been happening with Whois and so forth very closely over the last few years and we've been actively participating in the community process.

The APWG is a not for profit organization, we have an organization organized in the United States and we also have an organization chartered here in Spain in Barcelona as a scientific institution. It is basically been dedicated and working over the last about 15 years to deal with identity theft and cyber crime. It started out purely as phishing but we deal with all the related kinds of problems where phishing, malware, botnets, and basically what they're used for which is to steal money and impact people.

So the organization is dedicated to dealing with those issues and we do a lot of data sharing, and we run conferences where we present research and so forth and we publish metrics. Our membership kind of fall into some various categories broadly but they're all organizations that have to deal with these issues at both policy and operational levels. So some of our members are responsible for defending networks, you know, pieces of the Internet, infrastructure that they're responsible for running and keeping clean and they of course have to protect their employees and their users.

An example would be banks. Here in Spain, Caixabank, which is the largest retail bank in the country is a member of ours and they're on our Board. They have to protect their customers and they have to protect their network and their online banking and so forth. We also have universities and they're responsible for protecting all their departments and running their networks and associations of networks dedicated to security actually.

We have - and a lot of these organizations are also running CERTs, computer incident response teams. And we also have maybe what you'd call security companies, these are companies like Symantec, the company I work for, iThreat, these are companies that are responsible for helping other companies protect themselves in a lot of cases. They're gathering threat intelligence, figuring out what things to block so it doesn’t get to you and those kinds of things.

Of course security is something that’s outsourced very often because it requires specific expertise and specific resources to make happen. So all the places where we work and interact are probably using security resources or services provided by an APWG member most likely. Law enforcement does participate and academics and so forth, but basically a set of organizations, hundreds strong, who are dealing with these issues every day.

They all use, to some extent, registration information to understand what's going on. A typical example is one of our organizations is dealing with a phishing attack, it’s attacking somebody and it’s trying to take advantage of certain users. You have to figure out, for example, is that domain name owned by somebody who’s the phisher? Is it something we can maybe shut down without affecting people negatively? Or is that domain name owned by an innocent registrant and it’s been hacked into? The majority of phishing cases actually are dealing with these sites that are broken into.

I mean, you want to deal with that completely differently, you don't want to shut down that domain name because somebody’s using it and that somebody is not responsible for the bad action that’s taking place. A lot of what our members are doing also involves understanding what's going on in the wider Internet. It’s often the case when there’s one domain name that's being used for a crime it’s very likely that there are some others associated with it.

We have to understand maybe what those domains are by linking them in various ways, including what name servers they're using, what IP addresses they're using and registration data has been a key piece of that. Registration data that’s accurate is wonderful but usually it’s the nice people who provide accurate information. Inaccurate information is also very interesting to us because we can tell if somebody is kind of trustworthy by whether they're providing good information or not; that is one indicator among many you can look at.

And criminals, by the way, are pretty lousy sometimes at faking their information. You can find out a lot about it from looking at that registration data. So we've been using this information and describing how we use that information as members for a long time here at ICANN. And of course the situation has changed; a lot of that information is no longer available and there’s some good reasons why the situation has changed with GDPR.

We want to figure out ways to work within the law and to figure out ways where the data could still be available for allowable purposes and that’s part of the conversation. So GDPR does tell us some things that we can work with. Let’s start there. Some of the recitals in the GDPR are quite useful for us. The recitals are sections of the GDPR that kind of give examples and some guidance on what the law means.

GDPR recitals 47-50 are useful; those are the ones that say when you're doing balancing about how to work with the data and who has access to it, some of the cases in which that balancing is justified and can be considered are cases about cyber security including the uses of preventing fraud, ensuring network and information security, the ability to resist malicious or unlawful actions, the ability to report possible criminals acts to the authorities and so on. And these are the things that our members do every day. And by the way, working with law enforcement is actually something that a lot of our members do very closely.

I want to emphasize the fact that a lot of what happens on the Internet of course as far as security is actually done by the parties who are responsible for the networks and so forth. Law enforcement gets involved in an extremely miniscule number of abuse cases or criminal - activities that happen on the Internet and that's because law enforcement has certain resources available to it but no more. Investigation of course takes time and prosecution takes time after that. Most of the lifting is done by the people who are running those pieces of the Internet.

And a lot of the stuff that law enforcement does relies on security companies in various ways. Some of them receive data from security companies about what’s going on on the infrastructure of the Internet. There are also referrals for example, if you talk about most of the major botnet cases that have happened over the last 10 years, if you read those press releases from Interpol or Europol or the FBI, you’ll see a list of companies that helped them; those are the companies that provided them with data, maybe provided them with the initial leads and not for profits as well, absolutely.

So what we’re looking for is to work within the law obviously. But we’re also looking for an access method that is predictable and rationale and repeatable. The current situation right now is if you do want to make a request for the information, it is none of those things; it is not - it’s an unpredictable process where every registry or registrar will tell you something different, they’ll have their own process.

They may or may not give you the information. The timeframes may vary widely and it will be basically about one domain name at a time. If you ask, for example, we have this problem, here’s what it is, here’s the evidence, are there multiple domain names that this registrant has? They will not tell you in general.

We're looking for a way to say we want to be able to make queries through RDAP, justify why those are being made and have a plan that offers predictability but also is going to be providing accountability for all the parties involved. Our members realize that, you know, they are going to have new responsibilities in this new world under GDPR. And that is kind of the way we’re starting to think about how a program might work. GDPR again is the guide because it says, you must do things like pay attention to data minimization, you must pay attention to the storage and the security of the storage, you must not keep that information for absolutely longer than necessary. You know, so we want to bake those things into whatever process comes out of this.

This means that some parties may not qualify. I mean, we would like to see a program where our members can get accredited and that’s going to be not only a process where first you have to be a member but then you're going to have to go through some sort of an examination and you're going to have to prove that you can do these things and you understand your responsibilities and you know that you can be audited and those kinds of things should take place.

And, you know, there will probably be some parties that may not meet those standards anymore and will probably be very difficult for individual researchers to meet those bars. Very small companies might have to make some significant changes to the way they do things in order to meet those bars. Some universities are set up pretty well to do these things, they actually have departments dealing with cybercrime and - or computer science departments that deal with these kinds of things and they have really good processes for these kinds of things, but others might not. So if you're a member one of those departments might be easy or you might have to start doing some things very differently.

So the GDPR does envision ways that accreditation and codes of conduct can take place, but now we seem to be in a kind of terra incognita where these kinds of things haven't been instituted and put together before. So I think we’re all in the same boat, we have to learn how to do these things and do them together. So that’s kind of briefly what we would like to see at a high level.

Ayden Férdeline: Thanks very much for that, Greg. I’d just like to get a sense for how many questions we might have in the room if any? Okay, if there are two we’ll go with Stephanie then Milton, please.

Stephanie Perrin: I was actually going to suggest that we also have Patrik Fältström and Rod Rasmussen at the same time. How would you like to manage this? I mean, we have Rod scheduled in a bit later. Maybe this kind of discussion should be more of a…

((Crosstalk))

Stephanie Perrin: …a round robin right now with the three of you here if it suits your time tables? Because I appreciate that you're running between meetings to attend here so we really appreciate you coming.

Rod Rasmussen: Starting now would be fine with me to do that. I think the, you know, I was going to talk a little bit about, you know, application of standards which is a natural dovetail to what Greg was just talking about and Patrik, I think you were going to talk about some of the same things?

Patrik Fältström: Well, I was to talk a little bit more about how an individual organization that deal with this sort of security-related issues from my perspective how they operate so it's sort of slightly different but I'm - it's easy now for me to refer to what Greg said, etcetera.

Stephanie Perrin: Well let's give that a whirl and see how we do and maybe you could pause for questions part way through, how does that sound? How many have we got in the queue already?

((Crosstalk))

Stephanie Perrin: Oh well, we can wait, right?

Rod Rasmussen: I was going to be brief too and, you know, thanks, Stephanie, for being flexible. I did want to come by here and spend some time and it's been a busy day and a already busy week, which, you know, is only going to be longer. I'm Rod Rasmussen, and you may know me as SSAC Chair but I'm not wearing that hat in this conversation, I'm just wearing the hat of a person who ran a company who had to deal with working under standards and dealing with a lot of other companies, working with them. A concept has been put forward that standards and certifications are a way of helping create a playing field or a system where people and organizations who need to access data may do so.

My company actually moved cyber security data - non classified cyber security data for the federal government in the United States back and forth between government agencies and between government and private sector. So as you can imagine we had a fair amount of things that we had to deal with and we also worked for financial institutions and other entities with very sensitive data that we had access to in order to do our jobs, which was kind of more of the protection and outreach and things like that that Greg was just mentioning.

So part of, you know, any professional industry, as we probably are all familiar with in our jobs of some sort of certification and standards bodies that set the rules. There are several, and for information technology, cyber security, etcetera, cyber security itself is an amorphous term, it’s kind of a catch all for things meaning things that happen on the Internet that are bad or good or keeping them from being bad.

But there are standards around information security that are fairly mature and that a lot of people use for managing or for - as a guideline for managing their data and getting certifications and these tie into things like insurance, and ability to receive, you know, the data from - and the certifications for being able to handle particular types of data.

In the case where we’re talking about here, these come under some of the ISO 27000 standards probably and then there are also frameworks (ANISA) Europe and NIST in the United States of cyber security frameworks that a lot of the member companies like APWG have already applied into their environments especially if you think about the financial sector, they have a whole series of things beyond the IT information that they deal with on a regular basis just to be certified.

These all lend themselves to being able to - to provide evidence that you as an organization have qualified to meet some sort of threshold, typically around the way you protect your networks, the way you deal with data, protect your customer’s information etcetera. Those do not speak directly to the problem we have here, you know, with GDPR and transfer of data related to registry, you know, RDS data. But there is a transitive property there I think that is interesting to look at as far as being able to do that.

So I think there’s value in looking at these and trying to pick out what ones would apply. In fact as far as, you know, looking what APWG is doing, that would be part of any kind of examination of a member applying, right, do you meet these standards and things like that. That actually makes it far easier to understand what level of maturity they're at. So that's all well and good and being standard certified is - gets you to that level.

However, still the operational aspects of what you do that really matter. And there are plenty of cases, you know, credit card companies require you to be certified - and I’ve just spaced on the name of the certification, credit card certification is a particular one and we all know about credit card breaches happening all the time.

((Crosstalk))

Rod Rasmussen: PCI, thank you, PCI compliance. It’s already been a long week. And I said I wouldn’t screw up anymore acronyms today. The - you know, pretty much everybody who’s had a credit card breach is PCI compliant so the standards do not equal you have a magic, you know, force field around all of your stuff.

It’s really the way you apply them and that gets into I think where a regime like we're talking about with APWG or others are talking with some sort of certification process, you know, the - while the standards - meeting standards is a nice thing to be able to more easily certify somebody or to, you know, have somebody else providing audits that give you some more assurance, it’s really having a program that’s ongoing and is focused on the area that you care about as far as what data is being protected and how it’s being protected that matter at the end of the story.

So while I think it’s a good place to look for guideposts and potentially even implementing some standards that would be specific to the kinds of information transfer we're talking about, it’s not - it’s only part of the solution; it is not a full - meeting standard certification is only partway there and it really is an operational thing going forward that you need to look to. And that's both in industry sectors and then as far as some sort of compliance regime probably around that. That’s just personal opinion on that. So I’ll stop there. I think that you got the gist of what I was trying to get to. I don't know if you want to go onto Patrik or take questions.

Patrik Fältström: Was it a quick follow up that you wanted to - okay, then I do a quick follow up. So to continue, I wrote some notes here what I want to talk about and it’s actually - I wanted to start and I promise I wrote this before Rod talked and we have not synchronized…

Ayden Férdeline: I’m sorry, we just need to pause for a moment while the recording…

Patrik Fältström: Oh okay.

Ayden Férdeline: Something happens to the recording.

END