CR - Replacement of WHOIS
Thursday, March 15, 2012 – 13:00 to 14:00
ICANN - San Jose, Costa Rica

Patrik Fältström: … that we in the Security and Stability Advisory Committee came up with, which to some degree can be interpreted as replacement of the WHOIS protocol. So we will go through the report so that everyone in the room sort of understands on where we are, but that the goal is to try to have as much time as possible for discussions later on. Also, we try to push forward to the discussions because we take for granted that people actually are relatively up-to-date on where we are.

So the WHOIS - and I intentionally say WHOIS here - is important for the entire community - that is something we hear all over the place. But one of the things, which is the problem of course, is that people just say WHOIS just like I said on the first bullet here, when in reality we might talk about the protocol, the data base.

So people might mean many different kinds of things, and one of the findings that we try to emphasize in this last document that we wrote is that the fact that people are not care enough on the terminology - that by itself makes discussions so difficult, so it's probably a barrier to be able to move forward and solve any other problems.

So SSAC has produced a number of advisories related to WHOIS 27, 33, 51 etc., and this is 51 that we are talking about to try to sort of compile this. And then as a result of that report, ICANN Board tasked the staff to produce a road map to implement the recommendations. And here is sort of the road map that at the moment we have a draft road map posted for public comment on the 18th of February, and then we have the workshop, which we are currently in. So with that as a background, Jim, maybe you can dive more into the report.

Jim Galvin: Thank you, Patrik. So very quickly - the reasons why this document came into existence and SSAC chose to address this issue, just that WHOIS discussions… WHOIS is a term used very generically by many communities and in many contexts, and we think that this really is part of the problem with talking about WHOIS.

So we kind of took a step back and looked at offering a set of recommendations that we thought would help to ground the discussions by being able to split them apart into the multiple pieces that people always talk about, and thus separate the communities and the discussions, and we're very hopeful that this will help to progress these discussions.

Next slide. So there are essentially three elements that we are proposing. The first of course, is Domain Name Registration Data. This is actually the elements that are necessary for supporting the life cycle of a domain name, and I think that these are in and of themselves in some sense very well understood. There are even some issues here in these discussions, but they are defined by the EPP Protocol, they are defined by what is currently in contracts, in registrar contracts of ICANN, and in registry contracts with ICANN. So there's a fairly well-known quality to what actually is the data that is represented when you're talking about WHOIS.

The second thing is Domain Name Registration Data Access Protocol. In the technical context, WHOIS refers only to this - this is the actual Protocol and this would be the WHOIS Protocol as it's defined today and that's what is used for the display of Registration Data. And what we want to do is identify a separate term, separate from WHOIS, for referring to this so that we can then talk separately about a different kind of Access Protocol in order to support other requirements which are coming now into existence, most notably of course, IDNs.

The third element in this taxonomy in trying to structure these discussions is the Registration Data Directory Service, and this is where we can talk about the

requirements of what needs to happen when you need access to the data. So the WHOIS Protocol today would be the Access Protocol. You need a mechanism for simply retrieving the data. Once you have the data, you need to talk about who gets to see that, and under what circumstances they get to see it, and through what mechanism it might be that they get to see that. So it allows us to look at law enforcement as a separate community which has its needs in getting access to data. You have intellectual property communities which have their needs in getting access to data. You have the public which needs its access in getting access to data.

It would be in this taxonomy name that one would have discussions about privacy and proxy services, so what kinds of walls or access restrictions that you might put on the data, the requirements on what needs to be there, and next slide.

So some observations about the structure of these WHOIS discussions and these three taxonomy elements. There is actually no uniform data model today for what in fact is registration data and we simply make that observation here, so having a term in which to talk about is a good thing. There is, as I said, there are some sources that identify what is Registration Data, but it has not actually been collected and put together and given that name and that label from which you can derive all of the rest of your communications.

An important observation of course, and I think we all recognize this, the WHOIS Protocol as defined today, the Access Protocol that we all use today for looking at Registration Data, of course there's no standard for handling non-ASCII data.

The standard itself is defined really just to be ASCII based. There are a number of solutions out there. Different registries have done some things to provide some IDN access in particular to data, and in fact to be able to display data which is internationalized inside of that registration data, but again, not a standard. And of course today directory services do not satisfy other legitimate needs for access to those services. So as I described, you have different

communities of people, law enforcement, intellectual property, the public, there are always privacy in proxy concerns. There are no uniform guidelines for what those people - how they should get access to the data and what it looks like, and so we think that it's important if you separate out the discussion of what is the data, the protocol that you use to access it, then you can separately talk about the kinds of access that different groups of people might have and you can set those discussions off and have them separately for those groups.

Next slide - so we have three specific recommendations and very quickly, obviously recommendation one was to adopt the terminology that we're proposing. Second recommendation is in fact to explicitly evaluate and adopt a replacement access protocol, so to call out that particular action item and task. You hear a lot of rumors about that being necessary and needed, but let's make it an identified work item and find a way to make that happen.

And then of course the third thing is to develop a uniform and standard framework for accessing data so we can really begin to address the needs of these other legitimate communities, who have legitimate need and access… should have legitimate access to this data. We can look for solutions that meet their needs and still meet the needs of the rest of the community, and of course any other privacy concerns that come about having with registration data. I believe that's it, next slide.

Andrew Sullivan: I'm Andrew Sullivan and I am talking a little bit about what is going on in the IETF. I want to emphasize for the record, that I do not speak for the IETF or any working group there - these are my personal views about what is going on in the IETF because people in the IETF get very sensitive if you talk about what they think.

Next please. So we have in the IETF, we are attempting to promote a protocol to replace the WHOIS protocol. This is called WEIRDS, it stands for Worthwhile Extensible Internet Registry Data Service, because the IETF likes to have cute acronyms. This is broadly lined with the CRISP requirements which

were published a number of years ago in response to the last time we went through the WHOIS pain.

This is our C3707 - you can read that if you'd like to - but the idea here is that this is going to be REST-based for easy implement ability. So you all maybe know what REST is, but if you don't, it's basically a way of producing web services, so that you just query a URI and you get back an answer. This means that the entire structure can be built fairly easily.

Next please. We had a BOF in Taipei. A BOF is a Birds of Feather session and Birds of Feather sessions are designed to assess the community consensus about whether a working group is going to get setup. So the IETF procedure is that you have this BOF and then you decide whether people are going to work on this thing, and if they are going to work on the thing, then you create a working group, and if you're not sure then you don't create the working group.

There was very, very strong agreement in that environment to do work on internet numbers. So the other WHOIS consumers, of course, are the RIRs and the RIRs use WHOIS in order to publish network number blocks and all the rest of that. So the real work around WEIRDS got started because ARIN and actually other RIRs started trying to replace this and ARIN has a test beds standing. I think LACNIC does and RIPE has one. So there's more than one test bed in the RIR world. So there's a very strong agreement in the room about doing that, but there was a lot of push back about doing work on domain names.

Now part of this is because there weren't a lot of registries, or registrars in the room, but part of this is because we've got a history in the IETF of having done this before, right. There are already three replacements of WHOIS protocols, and none of them have been implemented and therefore, people are a little skeptical that this time it's going to work. There's no agreements right now on proposed charters, but it looked like we were making some progress on the WEIRDS mailing list and therefore, we're going to have another BOF in Paris.

It's very important to understand something about the IETF procedures - there is that so far as I know never yet violated rule that you can have a third BOF. You

cannot have one, so this is basically our last chance. If you don't get chartered to do this work now, there won't be a working group. Next please.

So this is the point I was just making, there's a maximum of two BOFs. There is a problem that we have right now, that there just hasn't been much work on names. ICANN has, because of the work of Francisco and Steve, has a test bed and has a draft and it's been active, but there hasn't been that much progress on names, unlike on numbers, right. We've got several test beds on numbers, we've got competing drafts, people are definitely converging on something and on the names there's a lot of question.

The previous work means that there are people around the IETF who think that ICANN will never get organized about this, and that the names community cannot come to an agreement, and therefore we shouldn't do any work on it. I would like to prove them wrong, so I would like people to come and participate on that, therefore next please.

There's this important thing that we need to understand. We have this slogan – "Rough consensus and running code." So the rough consensus is seeing that people are doing work on these things. We have comments on the mailing lists, we have evidence that people are in fact converging on a solution.

If you've ever managed any kind of project, you know that your employees come to you and they say, "I've got this really great idea and I have no idea if it's going to work," you're going to tell them, "Well, come back when you're ready, because I'm not giving you a bunch of my time in order to do that." But if they come and they have a plan for how you're going to get there, then that's the kind of thing that you're going to support.

Similarly in the IETF, the area directors - who are basically managers of the working groups - what they're going to say is, "You people seem to have an idea of where you're going therefore, I'm going to support your work. You people over here however; you're too fractious. I think that's going to cause me a lot of work, so I'm not going to agree to it." That's really what's going on there.

The other thing that is very, very valuable in the IETF is the running code part. So if we see evidence of people doing work on these things, we see evidence of test beds, we see those kind of things - that's the kind of project that a lot of people get excited around the IETF because they say, "We can make this work really well for a lot of people." You know, the IETF is good as sort of suggesting how interoperation is going to work but it's not very good at blue sky kinds of problems. Another important thing to remember - work is officially done on the mailing list. If you can't make it to Paris, it's not the end of the world, next please.

So here is the mailing list that lists there weirds@ietf.org; if you want to subscribe, please do. You can follow that URL and it would get you to all of the drafts that have ever been submitted about this topic. There's the subscription link there, but that's also where you can get to the archives.

I really want to emphasize something - it's very valuable to see reviews of drafts, to see people saying, "I have read this and here is a comment." That's much more valuable than hearing, "I plan to work on this in the future." Planning to work on it in the future is a nice sentiment, but it doesn't provide evidence that there is work being done, whereas the work being done is the real evidence to an area director that, "Hey those guys are going to do the work. They're not going to create problems for me," and I think that's all I had to say. Oh well I should say please do come in and participate, we would appreciate it.

Francisco Arias: Thank you Andrew. This is Francisco Arias from ICANN – I'm the gTLD Technical Liaison. I just wanted to add a little bit about the drafts that are Marty's drafts from the RIRs and there is one that we have written, but it's not just Steve and I - there is also something on the second version we had another two authors. One from my CO that are from CNNIC and so we are trying to make this bigger.

Now I'm going to talk a little about the draft roadmap. This is the direction of the Board, the Board requested staff to write a road map to implement SAC 051.

In the current environment what other things linked to the support that have happened in the past? There is GNSO WHOIS service requirement report that was published a few months ago, it placed by requirements for WEIRDS is very similar to what has been published elsewhere for example in RFC 3707.

There is also joint GNSO and SSAC internationalization registration of the working group and they just published their final report. There is a similar recommendation in this final report to replace WHOIS, and we also have the WHOIS Policy Review Team. This is part of the reviews described in the Affirmation of Commitments for ICANN, and they released a draft report which has two interesting recommendations related to this report.

They recommended to internationalize the WHOIS service and they also go beyond that and suggest inclusion of our requirement to offer this internationalized service in contracts for gTLD registries and registrars. This is of course in draft and it's still being discussed.

Regarding the roadmap that we published and is still up for public comment - what we are suggesting on the recommendation from SAC 051, which is related to the new terminology, we think this is a simple thing. We believe that we should provide a summary of the terminology recommended and share that with ICANN staff and other stakeholders through the support organizations and Advisory Committees' Chairpersons.

We are thinking that it should be a transitional period in which their preferred terminology in documents will be included, but also incorporating them with all terminologies so that people can get used to the new terminology and better understand what is being said in the report.

The other two recommendations from SSAC we believe can be summarized in replacing the WHOIS protocol and here we are proposing a multi-prong approach - first we think that ICANN should promote the participation of ccTLDs, gTLD registries and registrars in the development of the replacement protocol in WHOIS in the IETF. On the policy side of ICANN we believe that

it should probably be a GNSO PDP with participation from a CCNSO, SSAC, and ALAC to replace the WHOIS protocol.

We also think that since PDP might take some time to develop, it probably would be a good idea to start including a negotiable inclusion of provision to operate these new protocols in gTLD registries and registrars as appropriate. For example, for gTLD registries - the ones that are more known to me - they have contracts that expire from time-to-time, we have two or three every year, so during the negotiation of those contracts, we could have a negotiation to include these new protocols included in the contract.

And finally we believe that ICANN should promote the adoption of this replacement of WHOIS within ccTLDs. Within ccTLDs there is no policy mechanism that is not an option to do what PDP to require ccTLDs to do something like that, so the only mechanism in order to reach to ccTLDs is to promote the adoption of the replacement protocol.

This is a graph that tries to explain the timeline of the roadmap. As I mentioned before, we will have first the gTLDs that are waiting to adopt the provision to replace WHOIS and that could be stopped immediately, and then we will have the development of the protocol and in parallel with the promotion of ICANN and with ccTLDs, gTLDs, and registrars to participate in the IETF.

Then we believe that it could be the first deployment by the first adopters and on the two sides - the gTLDs and the ccTLDs space. Then at some point the PDP will act and conclude and hopefully that would say that gTLD registries and registrars will have to adopt the new protocol, then the remaining of the gTLD registries and registrars will deploy the new protocol.

So what are the next steps in terms of the roadmap - as I mentioned before this is still for public comment. The public comment ends on next Sunday; however by the new rules of the public comment at the start of this year, we will have a reply period that will end on 18th of April. We already have some interesting comments in the public forum.

There is the URL in case you are interested in submitting any public comment. Because of the importance of this topic, we are ambitioning that there will be a second version of the roadmap that will be published sometime between April and May, and so it will be a second public comment period. We think that we will be on time to finalize the roadmap so that it could be a Board and community action in time by the Prague meeting at the end of June. Thank you.

Jim Galvin: Thank you very much and we are on time - excellent. So with that I would like to open up for 30 or 40 minutes of discussion.

Michael Peddemors: Gentlemen, my name is Michael Peddemors. I am President, CEO of Linux.Magic and we do a lot of work with the WHOIS technologies in a lot of different areas. I guess you could say it's closely aligned to some of the requirements you may have from law enforcements.

One of the important things is that we use mass WHOIS queries and we have to be able to, of course, do large scale audits. We may have cases where, for instance, one of the common things is in spam protection where one person may be using thousands of domains and we need to look for similar patterns or similar owner for these domains, and what really concerns me is a couple of things.

Number one, of course, is insuring - and this is going through a lot of the different programs - insuring that the owner is always made public and freely accessible. As we say when a person is asking for a public IP or a public domain name, they should be able to at least put a public face on it. So this is one thing that we believe should be a strict requirement that gets observed through all three processes. And the other thing is that we're really concerned about some of the web-based protocols.

Although this may satisfy the general public who may like to find out who owns a specific domain, some of these web-based REST API-style protocols may not

be as easily incorporated in other technologies, or perform as well for large mass queries. Thanks.

Jim Galvin: So in fact, it was exactly those kinds of issues that prompted the production of SAC 051, because I think what's interesting is if you separate out the directory service itself, you've identified at least two communities of access that you want for that and things that you want to do for it, and our observation in that is that those don't need the same kind of solution. You can have multiple solutions to the problem of display of data.

So I mean Andrew will tell you about what's going on in the IETF. And yes that's a relatively focused point of view and a particular type of directory service based on a particular kind of protocol. But one of the reasons for separating out the discussions between the protocol and the service is you can have different protocols to support different services. We want to encourage the discussions to go in that direction as opposed to focusing on one protocol and one solution, right. One of the problems with WHOIS, as you say, WHOIS – they've only got one protocol available to you, one display mechanism and we're trying to build stuff on top of it.

Andrew Sullivan: The point of the protocol replacement proposal is in fact, that it would be easier to build those services rather than harder. The problem that you have today with Port 43 WHOIS is that it is essentially a bag of bits. You make a query and you get back a pile of stuff on the wire and you don't know what order it's going to be in, you don't know what data is going to be in there, you don't have any data model whatsoever.

So part of the goal of WEIRDS in fact, is to bring structure to all of that so that it's machine readable on output format and on top of that of course, you can build these other services. You don't like the thing that comes back from W-GET, and then you can produce yourself a different kind of thing.

It's a machine-to-machine protocol really that it just so happens will fit nicely into an HTTP type of response. Look, we could today replace the WHOIS protocol with IRIS. We already specified this thing we could do with our WHOIS. We've got lots of technologies for these things - they were all failures technically because nobody deployed them. We believe that part of the reason nobody deployed them is because all of the clients in the world had to be replaced and you needed a brand new server, and you needed to understand the transport.

Therefore we're going to understand the transport part by saying, "You all know how to use HTTP libraries - that's easy to do, any junior programmer can do that - there's a million libraries out there to do that, so that's very easy and everybody's got a web browser, so that's very easy."

So now we don't have the client or the server problem and if you're riding an automated machine like you are to receive those WHOIS queries, you can do that as a web client - pull the things back in. You don't have to render it of course; what you're going to do is process the underlying data instead.

Francisco Arias: You just said everything that I wanted to say about the machine readable code and everything.

Michael Peddemors: Sorry just to make it a little bit more clear, I've actually almost heard two opinions right in there while you were talking. As you mentioned, if there's multiple protocols that can access the same data, that's something I believe I think will be more accepted. The people who have the requirements for simple, clean, fast WHOIS can still use it, the people have other additional could be able to use say a new REST API model which may have different, whether it's performance factors, or things of this nature.

I would like to see in the recommendation a clear opinion when you say is this a replacement will we maintain the ability as we improve the data model, can we

not maintain both the standard WHOIS 43 protocol and look at this as an additional protocol?

Patrik Fältström: Francesco.

Francisco Arias: Thank you Patrik. I just wanted to clarify one of the things that you mentioned about the data that should be in the protocol I just want to clarify, or in the service tool let's say. This is only about the protocol, this is not about what data should be there - that's governed by the contracts that detailed these registries and registrars have, or by the local policy that the ccTLD has, so this is not about changing that. Just wanted to clarify that.

Andrew Sullivan: I did want to respond to that. Are we going to keep Port 43 WHOIS? God, I hope not, that protocol is broken. It is awful; it's been broken for more than 10 years, much longer than that really. It's designed for a different environment, we should get rid of it.

Patrik Fältström: Jim.

Jim Galvin: And just one quick point of clarity, it was not my intention to at all sound like I said something different than what Andrew said in my initial response. So to speak to the issue of multiple protocols or one protocol, I'm just suggesting the door is open for that discussion if it happens to work out that way, and other than that we'll see with the community and where engineering takes us.

Patrik Fältström: Okay Mark.

Mark Kosters: This is Mark Kosters. I'm the CTO of ARIN and we actually deal with numbers, but we also serve up a WHOIS service. And one of the things that we did just about a year ago is put out a RESTful API and essentially built upon that RESTful API WHOIS interface as well, so you can come off of Port 43, or Port 80, or soon Port 43, if you wanted SSL transfer.

What we have found through this is that people are moving from Port 43 to the RESTful interface, so much in fact that over 50% of our traffic is now going to Port 80 and soon Port 43. So we're seeing the consumers actually voting basically with traffic, what we're seeing, and we're seeing immediate use of people building on applications that use us, and it's really quite amazing. So we anticipate Port 43 going away.

Jordan Buchanan: Hi. I'm Jordan Buchanan just representing myself. So thinking about this and looking back, as you guys know, we tried this before and hasn't gone that well, and I was reflecting on that. Certainly one of the reasons why is that it requires big changes in clients and servers and that's annoying, and that's I think coupled with the fact like that is what are the benefits that we get out of the change.

I think so far they haven't been telling enough to justify people to get their systems and try to start over with something. It strikes me that we're not hitting the mark here either. I'm not at all convinced from seeing this and from having watched this is for a long time that this is going to be a better solution, that it's going to get more adoption than some of the stuff you've done in the past.

The point here is to support IDNs, great. It seems like a heavy weight change in order to accomplish that, but you know maybe Mark's right and people really want RESTful, that's great and we can switch to something that's RESTful and we're just trying to solve that problem, cool. That seems like a totally tractable sort of problem that we can solve in this face, and then we should say that's goal.

If the purpose is though, to try to like address this much broader set of issues in the domain community about what's WHOIS data, what it's being used for, who's allowed to use it and so on, it's way premature to try to solve those problems without any consensus from the community about what the solutions are supposed to look like, and we're nowhere near that consensus.

And I think more importantly it doesn't seem like there's concepts, I guess what I'll say is that we're replicating essentially the same system with any protocol, right? There's a pile of data, we're going to put some structure around it, great. We're going to have a new way to interact with it, great, but the fundamental problem that we have is that just being able to have a pile of data that everyone could get at doesn't seem to be sufficient in order to resolve everyone's needs.

So, you know you guys mentioned privacy briefly, but you know privacy comes up over, and over again as a really critical element of what's missing from the WHOIS debate today, so that should be a core requirement for trying to solve these things as opposed to a secondary one.

The other thing if you look at the WHOIS Review Team they mention the distinction between contact ability and data availability, and I think the technical community could be thinking about is there a way to address bridging that gap as opposed to just saying, "Okay here's a pile of data and we're going to give people access to it again." We keep doing that over and over again and it doesn't seem like a very interesting solution to the set of problems we have.

Patrik Fältström: There are lots of people that want to speak, so please be brief. Andrew.

Andrew Sullivan: Two responses to that, there are two things I want to draw out. The first thing is that the protocol is just a data access protocol, but by structuring the data, we actually could provide differential access to different people. That's actually part of the reason that we would like to do that, but who gets what data and so on - and that is from the protocol level's point of view, just a policy matter and

it's none of our business. So the goal is to provide those things, but we've already done that before, we've done these other things.

You're right about the first thing, and that is that there isn't - as far as those of us that are just propeller heads can tell - there isn't consensus in the ICANN community about whether it's really pressing to replace that protocol, and that is actually the reason for the skepticism within some people in the IETF against doing the names work right now because we don't know if it's going to succeed for exactly those reasons.

Patrik Fältström: Jim.

Jim Galvin: So you made the comment that this is about supporting IDNs and I want to clarify that statement just a little bit, because actually the problem is broader than that, and I think that we actually are taking a significant step forward here by doing what we're suggesting to do here.

This is about internationalization, which is more than just international domain names, okay, because it's also about the registration data itself, and what you need to do to it to be able to store it.

Now you could argue that the EPP protocol and what's going on in the registration process does actually support that to some extent. Certainly all the tools are there to make that happen, but there is no opportunity for a standardized and uniform way to display that data, nor to query data, which is internationalized and that's a significant issue.

Our big assertion here in trying to promote and progress a structure in taxonomy is to be able to separate these discussions. I think there's a really important and huge win in being able to pull out the data and decide what it is, talk about what it needs to be and then recognize that it does all have to be tagged with language and script so that that's there.

Having separately to talk about services and being able to build something which will then display that data and leaving things the way that they are is really important and I think that's a huge step forward. It gets us more in the community than we had before.

All the rest of the discussions that you're talking about - those have to be had. We do have to those policy discussions about who gets what access and when, and by what means, and this is not intended to solve those problems, but it's intended to scope them into a spot where they can be talked about and not disturb going forward with things that we really do need today.

Wilfried Woeber: Yeah, Wilfried Woeber from the numbers neck of the woods originally and on the RT4. A couple of observations here. First of all who will support your most recent statement about the stuff the whole system has to rather sooner than later support internationalized data? This is not just related or limited to IDN domain names. That has to be supported across the board.

The second observation is that being involved in the RIR business and in the numbers registry stuff, at first glance we always think that we don't have anything to do with the names people. That first impression is blatantly wrong because even in the numbers registry we are keeping and maintaining domain name registration data for the reverse delegation tree.

So we are definitely interested to sort of see something which is applicable to the names world and is as well applicable to the numbers world, because there is no good reason to try to do two different things for something which is very similar. Sort of just as a statement, as an observation.

The second thing observation and I have a long list of things, but I'll try to be brief, is the statement there was no implementation of the previous modernized WHOIS things - they're very implementations that the RIR agency was running a service for CRISP/ARIS stuff, and it was more or less complete, but the community sort of decided not to use it, and I don't know what the lengths of

the service by ability was. I think it was two years or something like that, give or take. My personal assessment is one of the reasons why this attempt failed was the fact that it was not speculate compatible with the installed phase, so I would not want to argue for keeping this real life perpetually, but it might be useful to think about mechanisms to make the existing installed base sort or the new stuff to make it more compatible with what we have.

The last thing is a question to everyone. Assuming that the IETF process fails in the sense that the IETF would not accommodate a new working group, what would be an alternative umbrella to do the work that needs to be done? I'm pretty sure that the work needs to be done whether the IETF agrees or not.

Patrik Fältström: Andrew.

Andrew Sullivan: Just on that last point, I'll be very short. The IETF is not saying that the work doesn't need to be done. What the IETF is saying - the skeptics who are pushing back - are saying, "We don't think it's going to get done. We think it's important and valuable and it ought to be done, but we don't think those guys are going to get it done," and that's a different problem.

Patrik Fältström: Jim.

Jim Galvin: I would prefer not to be thinking about back-up plans, because I think that it does need to happen and I think that there is a greater confluence now of activities and needs and requirements coming together than we've had in the past. I'm hopeful that that would lead to success this time around, but only time will tell, so we'll see.

Male: This is (Inaudilbe). I have one question for Andrew. As Andrew said that this WHOIS protocol [three times]. I know the ICANN team [is leading the investment] process. What is the probability of success in the current placement?

My second question to Francisco that what is the role of test bed part of the WHOIS protocol? You said that test bed is established by the WHOIS protocol. What are the main activities, and what is the role with that aspect?

Patrik Fältström: Francesco, this is for you.

Francisco Arias: I'm sorry I couldn't hear the question.

Male: I said that Andrew said that San Francisco is overlooking the aspect of the WHOIS protocol. What is the function of the test bed? What is the role in this replacement?

Francisco Arias: I think what Andrew mentioned was that he and I worked on a pilot implementation of WHOIS for names and that's what Andrew was mentioning before. We haven't put this public, and we probably should I guess in order to give people a look and feel what it could be, but that was what Andrew mentioned.

Patrik Fältström: Jim.

Jim Galvin: So to respond directly to your question of what is the probability of success this time. Rather than try to give you one, I'll make some observations for you. Things succeed when people participate, they insist that they want it and they use it, and that's really what Andrew has been referring here. I mean we've had three examples of replacements for the WHOIS protocol. They just never really got traction because nobody wanted to offer them and nobody would use them.

The only observation that I make is I think that there is a greater expressed need now for a replacement, notably because of internationalization, and I think that this community really has to drive that and really has to participate. We need more people in the WEIRDS Working Group, in the IETF and to do that you only have to be on the mailing list - doesn't require any other participation but that is the way that makes that work.

And also in the ICANN community here, there are opportunities for WHOIS discussions and WHOIS work - activities, various working groups doing various things. You have to get engaged, you have to talk, and you have to say that you need this. If internationalization matters to you, then you have to make this work happen and you have to make it successful.

Patrik Fältström: Francisco.

Francisco Arias: Yes, regarding the participation in the (Inaudible) and the IETF, I just wanted to mention that there has been already a few engineers from different name registries that have shown interest in participating, and actually have been involved in discussions. I'm talking about engineers from big name registries - Bay Sign Affiliates, NeuStar, ccTLDs, Outsource, CNNIC, .mx, Nominet, and I'm probably missing others registries over there and [Central NIC] I believe, too. So there is already a growing community that is trying to make this happen, but it would be great to have more participants in the IETF effort.

Patrik Fältström: I'm sorry I missed you.

Male: That's okay. I'd just like to come back to the issue of splitting the issues which are part of the whole game of making it more accessible to find a solution. One of the things that boiled down yesterday during the RT4 internal working group meeting was that the very fundamental question of what sort of data should be there, how the structure should be, and how it should be displayed.

That's actually independent from the transport protocol, and this is something that we are going to sort of most probably have some bearing in our RT4 final report. Sort of trying to make it obvious that this whole internationalization thing needs to be tackled whether it is in the traditional framework of old WHOIS Port 43 or whether it's in the framework of a replacement, or a developing thing. It doesn't make any difference - we have to deal with that very special problem.

And in the vicinity of that, I'd like to come back to one of the statements in one of the very early slides that there are different sort of, consumer needs or consumer expectations for the various constituents like law enforcement and private citizens.

What we found pretty early in our work is that it is not that clear cut, so there is not really a concise definition of what is law enforcement. That's one of the reasons why we ended up with the phrase like law enforcement and the security industry - that's probably going to be a phrase used in the final report. The reason is that in some countries law enforcement is really law enforcement in a traditional sense.

In other countries, actually the entities doing the day-to-day work for the sort of low hanging cases is actually not law enforcement, it's entities that are established as private entities, private companies. And then as this gentleman so properly pointed out, there is also the security gangs, and they are not necessarily law enforcement. In some countries they are, in some markets they

are, and some others, they are just anonymous help yourself groups. So it will be pretty difficult to come up with the concise definition of, this is law enforcement, this is domains, this is the man in the street, or the woman in the street, just as an observation. Thank you.

Benedict Addis: My name is Benedict. Can you hear me on this? I'm the representative of law enforcement and I think probably the only one left come Thursday. There is one - oh, there they are.

So we thought and mulled a lot of over this and I think while we haven't come to an answer, I think of at least defined the problem. So the issue is this - law enforcement have due process which is a way for legitimate reasons to get information that is normally private. So they can serve due process normally in their country by a relatively simple process to obtain information the company stores about its customers to prevent, or more normally to detect a crime after it's happened. So law enforcement uniquely has that structure.

So if we're talking about having a WHOIS that has an opt-out which seems to be a general feeling this week, then law enforcement only is going to have access to what was previously public information.

So the questions is, is there any way of phrasing some regulation around this to give access to WHOIS to people that should and ought to have it in order to keep the network security? I've got no answer to this tool and I am wondering - it's always been a perennial problem at how to recognize law enforcement and whether we should try to distinguish between law enforcement that is trusted, and law enforcement that is not trusted. Does that question even make sense?

Patrik Fältström: Andrew.

Andrew Sullivan: It's Andrew again for the transcript. Part of the difficulty that we have with the existing WHOIS service is that there is no query side authentication. That is, you have no idea who is sending the query, there is no way to tell. It's completely anonymous and it's uncontrolled.

There is nothing we can do, given the protocol that exists to fix that. Now I have personal views about what data ought to be public and not in the WHOIS, but as far as I'm concerned that's a policy matter that as a protocol geek that I don't care about. But as a protocol geek what I care about is providing the structure by which those differentiations could be made. And in fact by using something like WEIRDS or some of the previous answers that have already been standardized, we could, for instance, provide different kinds of data to different kinds of people.

Mail anti-abuse people - I can speak about that a little bit. Mail anti-abuse people, for instance, don't really care who owns the domain name. What they care about, right, is the registrant of name A the same as the registrant of name B, which is the same as registrant of name C - that is a different question than who is this guy. And by separating out those things and providing the technical mechanisms by which those things could be answered differentially, we increase the ability of the protocol to solve or to support these different policy answers.

So the goal of the protocol level development is actually to make those kinds of differential queries possible, whereas in the policy question, that's a different matter. Who gets to do this and all the rest of it - it's not my problem.

Benedict Addis: Thank you very much and that's very clear. I feel that's kind of a long way off and in law enforcement we're going to be resigning ourselves to manual queries, or at the God awful [MLAT] process for a long time to come.

Andrew Sullivan: Just to pick up on that, it may be a long way off from the policy point of view because there was a PDP on there and as far as I can tell, those are actually

infinite in length. But from the protocol point of view, Andy Newton, who has been working on this, sort of wrote down one day on the mailing list, "Well this is roughly what I think the names thing ought to do," and by the next afternoon, so 24 hours later, somebody posted to the list, "Well I implemented that one, here you go."

You know this is not final. The whole point of this was to illustrate that a RESTful system is to relatively easy to implement. So one of the complaints about IRIS the last time was, I need a 6-month budget of eight developers in order to develop IRIS servers. And what we have proven by example is that what you need is like four hours this afternoon and you can build something that will kind of work. It won't run an internet scale, but at least you can implement it. So part of the goal of this project is to solve that.

Male: I'm sorry, I just wanted to make a little fall-off because I'm also hearing a little of a cultural thing. You know we're talking about using your expression "propeller heads." Well when Port 43 and the WHOIS protocol had had some very specific design concepts around it, one was simply openness. What I'm talking about is that it was never built-in restrictions that who can access the WHOIS, data etc.

Now you're talking about, we have a lot of hosting providers are ready to jump onboard with a brand new protocol so that they can implement these new technologies, but I just want to point out that those original 10,000 propeller heads who had a reason for WHOIS, you know, when we're talking about new needs out there, well yes, there might be millions of new queries as the gentleman from ARIN mentioned, but these are a different type of people, a different type of need, and we can't forget that the original need of those original 10,000 propeller heads are served by the original protocol. And the biggest problem that we have with that protocol is simply that there hasn't been standards of how that data should be rendered, displayed, and I think that's a critical issue that we can talk back into on protocols.

Male: Now I put myself in the queue. To some degree where I agree with you is the protocol itself does not include any authentication mechanism that can be used for authorization proxy data, but there were absolutely the ability for people to know that what was available over WHOIS is something, just like you say, that was available for everyone. So there is still the case that you have two levels of access to the data - either it's available in WHOIS, or it's not. So it's not the case that everything was available.

Patrik Fältström: Andrew.

Andrew Sullivan: Just one more thing on the history. Historically it's simply false that the WHOIS was designed to do what we use it for today. That is a myth and it's a pernicious one. The WHOIS was created for the ARPANET. We knew who everybody was, the list of people was published on a piece of paper that was mimeographed and sent around to every site, and you are talking about the site managers, not the people on there, right? The site manager was like MIT. It's a guy at MIT because he's running a network. It's an experimental network and you know who everybody is.

That is not the condition of the network today and the idea that we have this protocol, and we've had it, and therefore we have to inherit it and continue to use it is simply false.

Arturo Servin: Arturo Servin from LACNIC. Also we are WHOIS operators. It's a nightmare to maintain that code. If you, right now you want to be a registrant, or another RIR for example, or an NRR, it would be a nightmare for them to run a WHOIS because the code, I don't know what is it, is very old. Our experience after we published the draft, we build a prototype in less than a month, and I'm a very

lousy programmer and I built an API in a weekend. So we really APIs and the HTTP protocol I think is the way.

Also, it let us build other kinds of things like applications, sort out authentication by the privacy, because we have the protocol that can do that with certificates. We can verify who is accessing the data and to give them special access to other kinds of data if the policy allows that. And also we can do the regulation that right now is awful difficult inputting, so I just wanted to point out what Andy said that. And the other people that build in the prototype in just one week was a [pilot] from APNIC that we just make and implement a draft in just one day.

Patrik Fältström: I think we are coming close to the top of the hour, so I would like to give the panelists an ability to give you last statements, but I see that there is quite a lot of interest here and I hope that the interest implies that all of you that have been speaking actually are members of this mailing list where the discussion is really going on and with that should we start with you, Jim?

Jim Galvin: Sure, thank you. I welcome everyone here and the participation we had here, and as Patrik said, and my closing comment is please, to encourage you to continue participate, find your way onto the mailing list.

I'm here representing the SAC 051 document, so speaking to that - we're trying to propose this new terminology and structure as a way to separate discussions. One of the fundamental problems that we think we have whenever we talk about WHOIS, is you get a group of people in the room, they have different requirements, they're talking about different things, and I would strongly urge people to look at this document, look at the terminology, think about it and either start to use it or come back and talk to use if you think we need something different, so that we can fix it. I think it's important and essential to making forward progress. Thanks.

Patrik Fältström: Andrew.

Andrew Sullivan: Thank you. I want to thank everybody for listening to the ranting geek up here today because I appreciate it. I believe that one of the things that caused us difficulty in the past was that the protocol and policy communities have not talked to each other enough.

So I very, very strongly encourage people who have needs – you have the needs - I got the protocol, but you have the needs - to come and make sure that we don't overlook at it again, right, because if we start building this stuff and we get it wrong again, it's going to be even harder the next time around to get it right.

Patrik Fältström: Francisco.

Francisco Arias: Thank you, Patrik. I just wanted to reiterate about the protocol ADIs to have a protocol that enables policy decisions, not a protocol that dictates the policy decision. So we could have a protocol that enables the different policy makers, depending on if you are ccTLD or gTLD, or whatever, you will have the policy that you want to implement. And I will also reiterate the invitation to participate in the IETF port and to provide comment in the public forum for the roadmap. Thank you.

Patrik Fältström: And with that, I would like to thank everyone for coming to this session and look forward to future discussions. Thank you very much.

[End of Transcript]