EN

BUENOS AIRES – Improving WHOIS: An Update
Thursday, November 21, 2013 – 09:00 to 10:30
ICANN – Buenos Aires, Argentina

MARGIE MILAM: This session is Improving WHOIS: An Update. We have a presentation here for you, from all the various presenters, to talk to you about different aspects of the ICANN program to improve the current WHOIS system. Many of you may have sat on previous sessions during the week about changes to the WHOIS system going forward, but this session is dedicated purely to the existing WHOIS system and what steps are being followed to try and improve it.

We have an action-packed list of items here. We will reserve time for questions at the end, and we'll cover various aspects of the WHOIS program starting with an introduction, all the way through various departments at ICANN and the work that's underway to try to improve the accuracy and efficiency of the WHOIS system. As I set the stage for this discussion, I wanted to provide you with a background of how we've gotten to this point.

The ICANN board, last year, took a look at the recommendations from the WHOIS Review Team, that listed a number of items that should be adopted by ICANN to improve the efficiency of the WHOIS system. AS part of that, the ICANN Board adopted an action plan that had a very detailed list of items that staff was told to implement in order to try to improve the WHOIS system in accordance with the recommendations that came from the WHOIS Review Team.

Note: The following is the output resulting from transcribing an audio file into a word/text document. Although the transcription is largely accurate, in some cases may be incomplete or inaccurate due to inaudible passages and grammatical corrections. It is posted as an aid to the original audio file, but should not be treated as an authoritative record.

That action plan was adopted last November. Since November there's been a series of activities that will provide more detail on today. It's summarized in an annual report that was published just before the BA meeting, and I've provided a link both to the action plan and to the annual report, so you can get a sense for all of the various activities underway.

With respect to the Board resolution last November, it really took a focus on trying to enforce the current obligations as they stand, and increase outreach and communications. That was one of the key aspects of the recommendations that came out from the WHOIS Review Team. With the action plan that was adopted, we had specific steps that staff agreed to undertake in order to improve all of these aspects, including communications and outreach.

It went as far as addressing some of the contractual issues that would improve the way that the WHOIS program was implemented. So since the last year you've seen a lot of activities, primarily related to the negotiations of the various contracts. As many of you know, we now have a new 2013 RAA that has significantly improved obligations that relate to WHOIS. We also have the new RAs for the gTLD program that also have WHOIS related obligations.

Then there's a series of compliance activities that Maguy will talk to you about in a moment, about how some of that will be implemented. With respect to the communications and outreach, a big part of it was to try to educate registrants and consumers regarding the WHOIS and how it can be accessed, and what the policy relates to. That brought in a series of improvements that relate to online tools.

I'll hand this part of the presentation off to Steve Allison, who will talk to you about the new WHOIS website, which is one of the key aspects of the activities to try to improve WHOIS.

STEVE ALLISON: Thank you Margie.

MARGIE MILAM: We're trying to set up the live… Is it working? Do you want to open a browser and project through Adobe Connect, and then we'll go back to the slides afterwards? While he pulls it up I'll give you a little background of the theory behind the website. What we tried to do was take a look at the WHOIS from the perspective from someone who knows nothing about the industry or what WHOIS is.

We tried to come up with content that would be very user-friendly and describe all of the obligations at a high-level. So as you browse through the website and Steve walks you through some of the features, keep in mind that it wasn't meant to be a description of legalities or contract terms; it was more a high-level description that could help someone who's new to the issue, to really understand and research issues related to WHOIS. Steve, I'll pass it onto you now.

STEVE ALLISON: Thank you. I'll just give a brief overview of some of the content that we do have in the site. After that I'll give you a background on some of the next steps that we have. We've put the WHOIS site into the six UN languages, and there's roughly 200 or so pages and links to other

resources that lead off of the site in each of the languages. Off the main page we have 'About WHOIS' section, and it really covers historical and background information on how WHOIS has come about, where it's at today, and some of the ongoing efforts to improve WHOIS going forward.

As well as this, it gives an introduction on some of the technical details that make up WHOIS and how it's architected. We envision that some of this section of the site will grow as we develop later phases of the WHOIS implementation, to include user guides and more detailed technical documentation around the WHOIS Lookup service itself.

We also have a glossary of terms that people see when they're interacting with WHOIS, as well as an 'About This Site' section that gives people updates on where we're progressing in the phases that we have for WHOIS. Lastly we have 'Submit Feedback', which is really the mechanism we've put in place to allow people to provide commentary not only on the content they see on the site, but also the functionality as well.

So I'd really encourage everybody to spend some time and feel that they really own this site, and think really critically about the content that we have in place, and suggest improvements so that we can continue to grow and make more accurate the presentation of it. Secondly, we have a 'Policy' section of the site. It's broken down into two main areas; one being implementation and the other, current issues.

Implementation would cover all of the work that's been done or is currently being done, as it relates to implementing WHOIS. Everything from the policies that we have in place, some procedures that have

maybe come from policy, or a list of all the registry agreements that have WHOIS-related information and links to those specific sections. Ongoing updates to registrar advisories and some of the benefits and responsibilities information as well.

Separately we have a 'Current Issues' section, which is really discussing ongoing issues related to WHOIS and an introduction to those issues. Thirdly we have a 'Get Involved' section, which really give an overview to site visitors as to what the mechanisms are for them to become involved at various levels of WHOIS.

Then a lot of the content in here will link off of the site to various Working Groups or links to mailing lists, or links to our meetings website so people can get a one-stop-shop for finding ways to become more involved. Fourth we have an 'Improving WHOIS' section, which as it is today has links to both the WHOIS complaint forms. We envision this section evolving over time with the later phases of the project, to include the WHOIS accuracy studies.

As we devote those phases, this section… The structure of how we present this information may change, but this is where we envision a lot of that information going. I'd also like to point out that I understand there are additional WHOIS complaint forms that may be developed over time and will continue to work to keep those synchronized with the site as well. Then last is the 'WHOIS knowledge center'.

This section of the site is dedicated to consolidating all of the information that we have that we thought was most relevant to the WHOIS discussion. It's organized into various categories that users can filter on or search directly against. Then it's organized as well by date,

so there's a chronology as well that a user can get. For each of those items, whether it be a link off of the site or actually a document that's located on the site, the user can get a brief description of that content and then pull it up as well.

That's an overview of where we're at, and just to give some commentary on where we'll be going with the site: we have two upcoming phases of the WHOIS website. The second phase being a Global WHOIS Lookup, that's currently underway. We've drafted an implementation plan and have also begun some of the development and design of that work. We're well underway for that. I believe we have two slides for this as well.

MARGIE MILAM: Before we move back to the slides, we've labeled this a Beta site because we really want to get feedback from the community on whether the content satisfies some of the needs of the community, and whether any of you have feedback on what we can do to improve, from a content perspective, the website, so that it becomes a central repository for any information related to WHOIS.

Steve, where can they send information if they have suggestions on how to improve the website?

STEVE ALLISON: Directly on the WHOIS website, under 'About this site' there's a 'Submit feedback' link that goes directly to our ticketing system. I think that's probably the easiest way to submit information so that we can track on

it. Okay, so Phase II is the WHOIS Search Lookup tool. It's currently being developed and we have an implementation in place.

We expect that we'll probably release that tool in two phases; the first of that phase being a WHOIS API with a UI on the website so that users can do their WHOIS lookups. It would be a domain lookup first and most likely return just the raw record back from the WHOIS query. Secondly, we're working to standardize the results and make the user experience a little easier to consume the information. So there will be a second iteration of that development that will standardize the output of that WHOIS lookup.

Timeline for that would be some time in January being the first iteration of that, and then some time in April for the second iteration. For Phase III of the WHOIS website we're doing WHOIS accuracy sampling and a reporting system. The approach we're currently taking is that we're going to develop a draft design of experiments to put out in front of the community, that will allow some commentary on the approach that we're going to take, or would like to take, for doing the WHOIS accuracy.

Then we'll get some feedback on that; on the messaging it presents as well as the technical methodology that we'll be going about. Once we receive feedback on that we'll be able to progress forward to actually doing a first wave of implementation on that.

MARGIE MILAM: Thank you Steve. Now we'll move to the WEIRDS presentation, so if you can pull up those slides? We have Olaf from the IETF to talk to you about that.

OLAF KOLKMAN: Good. I'm Olaf Kolkman from the IETF. I work for NLnet Labs and it's a habit in the IETF to have people not speak for the IETF but about the IETF. I do have a function in the IETF as Working Group Chair for the WEIRDS Working Group. I do that together with [Murray? 00:16:15]. He's provided some of the slide material.

WHOIS – WHOIS, as you saw on the website demonstrated a minute ago is not an acronym, it's actually the joining of two words; 'who' and 'is', as a basic question. What is it today? Well, from a technical protocol perspective it is an incredibly simple thing. It is ask a question and you receive a free form answer. Type in one line, press enter and something will return. Open a connection, type in a line, get a return, close connection. That's the technical description.

This is probably one of the easiest protocols in the IETF that there is. It has received an update, or two maybe. RC 3912 is the latest specification. Let's see if something drops from this page. It does. This is an example. You type in a command line saying who is… with some question around it, and you get an answer back. What you see here is that the question that you ask is actually… The structure of that question is context dependent.

It differs from the name registries and the number registries. The WHOIS protocol is a protocol to access information about registry data for numbers as well as for names. By numbers we mean IP addresses and so-called AS numbers. How you type in the address is context dependent, and what you get back is also context dependent, and how it's returned is dependent on the server you talk to.

That gives you a bunch of shortcomings. There's no differential service for WHOIS. It doesn't really matter who asks; law enforcement or the general public – you get the same information back. That might have privacy issues, or some people want to see a little more information than is disclosed to the general public. The protocol as such doesn't tailor for that.

It's not internationalized. Things like my name in Chinese characters is something that is very hard to do in the WHOIS. The response syntax is unstructured, which means that you cannot have machines interpreting the results, and automation is difficult. That is a very big reason to lose standardization, because you want to do automation. There's not a standard re-direction capability. If you end up at the wrong WHOIS service to ask a question, there's no standard way to get to the next one.

Some work has been done in the past to improve on that, and specifically with IP addresses there is something that works there. There's no standardized way of asking where to go if you have a specific question. What is the WHOIS server that I have to talk to to get an answer? I already talked a little bit about the what and the how. The how is the method of the data delivery, the what is what data is actually being delivered.

In this Working Group we try to talk about the how and not about the what. Obviously there's an interaction between the two, but the what is really a policy question that's to be answered by the community. That depends on what community you talk to; ccTLD community, local community, ICANN community, RIR community. It's a local question. It's a policy question.

WEIRDS stands for Web Extendable Internet Registration Data Service. The word WHOIS doesn't turn up in that acronym, and that's important. We talk about registration data access protocols. That is to provide a little bit of clarity and distinction between the what and the how. Registration data access protocol is really talking about the how.

The Working Group is chartered to deliver a protocol that's supposed to solve all these issues that we just talked about – internationalization, structured reply syntax, the issue of redirection and authentication for different service is being tackled.

One of the more important features is that we try to do this over HTTP using a fairly standardized architecture called RESTful. RESTful, I'm not going to explain this slide in much detail, but it's basically web application services. It's the way that WWW works. It's basically the web architecture that we try to utilize. The benefit of that is that there's a whole slew of libraries and programming interfaces available for people who want to implement this protocol.

Earlier trials to improve on the WHOIS service, such as IRIS for instance, had the problem that they were very complex from a protocol perspective, and very hard to implement. That mistake we try not to make in this case. As a result, what you will see in the future data registration access protocol is that you can basically use queries that have structures, like URLs that you would type into web browsers and that are fairly easy to parse for humans, and very easy to parse for machines.

IETF works with charters, so there is a plan of action, a description of what needs to be delivered, what is in scope for the protocol and what is

not – or in scope for the work and what is not. That usually comes with a set of milestones or plans. The current plan is depicted on this board. There's a bunch of drafts that will become RCs at some point. Some of the work has been done – the framework of using HTTP, the security document has been done.

We're currently working on what needs to go in a query and what needs to go in a response, and how to structure that response. Then there are three other protocols that we will need to complete. The dates that you see there are the current plans. We'd previously had more aggressive milestones but didn't get the problems that needed to be solved, solved, to get to a last call in the Working Group.

That doesn't mean that there will be a final specification at that date. There still needs to be an IETF consensus call, which might raise some questions that need to be addressed. [We'll never know? 00:25:08] in the last call. But at least the Working Group has a fair sense of things being done once you've reached a last call within the Working Group.

What are the major issues currently that are being discussed in the Working Group? Search, how to do search. There needs to be a search capability. It can be very extended with all kinds of patterns that people might want to type in – matches on the middle of names or matches on the middle of domain names or entity names. People might want to do regular expressions or all kinds of other stuff.

What the Working Group chose to do – and it's a general approach to get a success on a reasonable time scale – is to keep things simple and look at what is being done today. What's important is that the Working Group name even has the word 'extendable' in it, which means that

even if we decide to do simple things now, it doesn't preclude more complicated things to be done later.

With respect to search, the Working Group seems to converge to a fairly simple method of search, basically based on a pre-fixed match, where there are some internationalization issues that are being taken into account, and those also lead to possible dealing with variants. There are some words in the document that describe how to deal with that – basically local policy.

Another issue is the bootstrapping mechanism – how do you find the appropriate server? There are a few thoughts on how to do that. Mechanisms might be slightly different for the names and numbers than for the domain names. There are three proposals currently on the table. I looked at this post-Vancouver meeting, which was two weeks ago, and at that time there didn't seem to be an emerging consensus yet.

Various ways to do this, and there's still active debate on what to do and how to do it. It might be that the three matters will just be described and see what gels out in practice. There's a bunch of minor issues. The redirect issue is one that we probably can have a better sense of how to do when we've tackled the bootstrap issue. The bootstrap issue is where to find a server initially, and the redirect issue is once you have that server, how do you get to the next server with more detailed information.

As an example, to not make this too abstract, think that you have a domain name in .uk and you want to know information about something in blaffle.co.uk. First you have to find the registration data server for .uk. That is the central place where you have to go. Once you're at .uk

then there might be a redirecting service to .co.uk, which has more detailed information. It's following that hierarchy.

We're working on object inventory. That's basically describing what objects are currently used in the wild to give us a good sense on what currently is in the document on objects, is actually covering most of the space. There is a description or a little debate on how to do proper RESTful – this is web architecture, so the web architecture people are looking at this and going, "You haven't specified this quite correctly."

I did see that a new draft appeared on the 19th of November and I believe that this issue is now… At least there is a solid proposal on the table. Some good news – well, the rest was good news as well, so this is not contrary but it's even better news, I should say. In the IETF we talk about running code and rough consensus. The rough consensus is only there when we reach the end of the practice, so we cannot put a tick mark on the deadline item, but the running code is actually there.

This is a Working Group that has running code on several implementations actually, while the spec is being developed. Usually that makes for a very good spec, and over the last few IETFs there have been inter-op test events prior to the IETF and the results of that have been fed back into the IETF and actually made the specs clearer and less ambiguous.

Coordinates, if you want to track some of this, if you want to get a feel of the work that's being done, here are all the URLs for the inter-web. If you want to know what's been done specifically during the last IETF, there's a link to the data tracker, where you can find the minutes of that meeting. With that, I'm done.

MARGIE MILAM: Thank you very much. Let's pull up the other set of slides and we'll have Mike Zupke talk to you about some of the changes related to the 2013 RAA and the implementation efforts underway.

MIKE ZUPKE: All right, thank you. Are we good to go? Okay. What happened? All right, is it all right if I drive these? Okay. My name is Mark Zupke. I'm the Director of Registrar Programs at ICANN, which means I've been responsible for implementing the roll out of the new RAA that was approved by our Board back in June. She's asked me to give a little background on that process we went through, what it means to WHOIS and where we are today. That's what I'm going to do.

The RAA really was amended in 2009 and pretty shortly thereafter there was call from people in the community who said, "That's a great start, there's still a lot of things we'd like to see improved in the DNS space, in a way of the RAA." Of course, registrars have the closest relationship to customers and they have the most available data about how domain names are being used.

There was a lot of sensitivity by people in the community, in law enforcement in particular, and I think similarly in line with intellectual property interests. This process really kicked off for staff after the gNSO had done some work for staff, in terms of negotiations with a document from the law enforcement community that had 12 recommendations for improvements. We took that, we took feedback from the gNSO and ALAC, and we had quite a bit of public comment.

We negotiated with the registrars for… 18 months is a commonly repeated number, but I think by the time we got it done it was probably at least two years. Eventually we ended up with a new RAA. This Agreement not only is different in substance but it's also a new form, so it includes a number of specifications, which all have the same full legal weight as the contract itself. I'll go through some of the more critical ones to WHOIS in just a second.

A few other points worth noting about the RAA – it includes a new amendment process, which continues to have deference to the consensus policy development process, but it allows, for example, registrars to come to ICANN and say, "We think there's a way to improve this Agreement.

"We'd like to see an amendment that allows ICANN to go to registrars and say 'we think there's a small issue that needs to be tackled.' We'd like to see if we can reach consensus and go through some public comment periods." But it wouldn't be the sort of thing that would require PDP – these are more like amendments in form. So that's something new. The other thing that's new about this is registrars have now got a right to renewal.

This means between these two new components we don't have to have different versions of the RAA going forward. We'll eventually have every registrar on what we call a 2013 RAA, but it will be amended in real time, instead of having this constipated process of every five years overhauling the entire agreement. So I think that's really a big win for the community. It's good for registrars, it's good for ICANN and it's good

for the Agreement itself. I think we can have a good, healthy Agreement going forward.

I mentioned the Agreement has specifications and some other provisions that are related to WHOIS, which are new from the 2009 Form of Agreement. New is a WHOIS accuracy program, which I'll go into a little more detail about. There's also a Service Level Agreement for WHOIS, which I've got another slide on. Registrars will be required to include, in their WHOIS output, an abuse point of contact. There is a specification on privacy and proxy services. I've got a slide on that so I won't go into too much detail.

Registrars will be required to provide access to WHOIS through IPv6. One change that was requested by registrars was the elimination of Port 43 WHOIS by registrars, if it's already being provided by a Thick registry. We will be looking in the future to Thick registries as the authoritative WHOIS, where registrars will not be required to provide that. Also requested by registrars was the elimination of the requirement to provide bulk access to WHOIS data.

That hasn't been eliminated wholesale, but pending potential changes in the marketplace that could be reinstated. Finally, there's been a policy in place since 2008 that was developed by the gNSO, which allows the reconciliation of ICANN contracts and WHOIS policies with a registrar's applicable law. So what it does is describe how a registrar can come to ICANN if it's received a notice that its WHOIS service is in or is potentially in violation of their local law.

The new form of Agreement recognizes that process in place and that some reforms need to be made to it. So there's a placeholder for

improving that and bringing it into line with some other parts of the new Agreement that recognize local law for registrars. I mentioned there's a WHOIS accuracy program specification. This phases in in different steps. The idea was, let's start at the easiest things, let's eliminate those problems and let's work towards the harder and harder things over time.

The very first thing that registrars have to do is they have to make sure that both the registrant and the account holder have record for a domain name, and obviously I would tell my gTLDs that there's no blank fields in their WHOIS data. That's pretty easy to do – there wasn't a lot of objection to that. The next step was to look at the fields and make sure that they're formatted appropriately as fields, so the email address looks like an email address, telephone number conforms to various conventions. It's the same with postal addresses.

Another one not particularly difficult to do, although maybe that's easier said than done. Then, as we're getting into a little bit more challenging work we've got a requirement that requires registrars to affirmatively verify the email address or telephone number, with both the registry name holder and the account holder, if they're different.

This means that a registrar must, within 15 days of a registration, contact the customer by either email, telephone or SMS, and require that person to respond to the registrar with some sort of a code – they could be clicking on a unique URL or something. It needs to be an affirmative response sent by one of those channels. That's obviously a very new introduction, or a pretty significant difference in the way that registrars are doing business now.

Finally, what I think we're probably looking at as the most complicated step of this right now, is that registrars will be required to validate addresses across fields. In other words, this house number exists on this street, this street exists in this city, this city exists in this province or state, the postal code matches… This is to the extent that it's technically feasible and commercial feasible.

The reason why this is not an immediate requirement is that there's some work now that staff and registrars, through a registrar Working Group must undertake to identify what sort of tools might be available to enable registrars to do this. When we come to do a rough agreement that these tools are available, then registrars will be given six months to implement and they will have to begin doing this address validation.

I won't mention the last point here, and that's the RAA. Steve's mentioned that in his slides a little bit. In addition there's a WHOIS Service Level Agreement. This deals with formatting and also responsiveness and timeliness of updates to the registrar's WHOIS output. While I mentioned that registrars will no longer be required to provide Port 43 data for domains in Thick registries, they will continue to have to provide WHOIS output on their websites.

To the extent that we continue to have Thin registries, or course registrars will still have their Port 43 obligations for those names. This describes a little bit about what is required of registrars. It's very similar to what is being required of registries, except they've got a little bit more in a way of round-trip time for queries, in response to queries.

Finally I mentioned that there's also a Privacy WHOIS Proxy Service Specification. We spend a lot of time with registrars talking about this

one and we had registrars and a negotiating team who we could have said, "Here's our list of things that people in the community have asked for. This would be the gold-standard; the ideal," and many of them could have said, "Yes, I'll accept that because my proxy service is already doing that."

I hand it to the registrars, they were very principled in this and they said, "We really believe this is more appropriate for a broader community discussion," so that's really where we ended up. Some of the really aspirational things are headed now to the gNSO for consideration. What we've got is a temporary or interim specification that will require registrars to be providing good disclosure about what their services are, what some of the benefits and potential risks are to their customers.

Their customers need to know in what circumstances their data, which was not previously public, might be made public, or in which circumstances the proxy service will relay communications to the customer, if it's submitted to the WHOIS contact. Finally, one of the things that was really important to staff in this is that proxy services run by registrars or affiliated with the registrars through some ownership relationship, will be required to Escrow their customers' data that's not currently in WHOIS.

In the 2009 Agreement we reached a compromised position, which allowed registrars or their proxy services to either Escrow the data or conspicuously inform their customers they would not do so. That option to 'conspicuously notify' their customers no longer exists. It would be mandatory for all registrars who are operating or are affiliated with the

proxy service, which is something that we on staff have felt strongly about for a long time.

This specification has some set clause. It will expire January 2017 unless the registrars and ICANN come to some sort of agreement that it can be extended a little longer. The idea is that we really wanted the gNSO to undertake the efforts to consider this and make whatever policy improvements that need to be made and not rely on this. I think ICANN and registrars, and everybody in the room, felt like there is opportunity in improving this, but the appropriate form was in the policy-making body.

Finally, just a little bit of data about what we've been doing to implement this new RAA. I may have mentioned there's a transition indemnity to the RAA, which allows certain obligations to become effective immediately and certain obligations to become effective January 1st 2014. It's become a lot less material the closer we get to next year, and so when we first rolled this out that was a topic of great concern for registrars and I think as we get closer to January 1st it's a moot point.

There are a lot of obligations that registrars who've signed the new Agreement are not yet obligated to undertake until January 1st. So to go through the nuance of that would be an entire session and frankly I usually spend two days going through the introduction to the new RAA with registrars so this an extremely condensed version. Just a few of the highlights – we've had 106 registrars already sign the new RAA. We've got 75 who've volunteered to sign it before their old Agreement expires, and so they're in the queue.

Most of them I think have submitted it within the last week or two, so Maguy, who you'll hear from soon, her team is concurrently reviewing them to make sure that they're all eligible for signing the new Agreement. We've been doing considerable outreach on this to make sure that registrars not only sign this but make sure they understand what they're signing; that they're getting ready for actual compliance and they're ready to go.

This is something that the registrars requested. This is something we heard from law enforcement. They said, "This is great, we have a new Agreement but we need some outreach and training." So staff has conducted… As of tomorrow we'll have conducted four outreach sessions. Three of them were two-day sessions. We did Los Angeles, we did one in China, we did one in Berlin a month or so ago and we're doing on here in Buenos Aires tomorrow, which will be slightly condensed. That will be a one-day session.

My colleagues, [Kate? 00:45:08] is our team's webinar guru. She's done a series of webinars and in general, every other week, at the same time, we do a webinar about different issues in the new RAA that are of concern to registrars; that are of the most concern. We set up a ticketing system where we can take questions from the registrars about the new Agreement. She's been taking those inquiries and building an FAQ document.

Also we've been using this to shape what webinars will be next. I think there are some good things planned in the very near future. So we've been really active in trying to make this a successful roll out. We want registrars on the new Agreement. Of course we want… The launch of

new gTLDs is dependent on this so we want that to be successful. We want registrars to understand what to do and what we don't want is Maguy and her Compliance Team to do a lot of work in cleaning up, because we didn't do enough outreach.

So we're really… I think we're doing an unprecedented level of outreach and training on this. The 2009 RAA we have a registrar training program for that, which I think went live some time in 2010 or 2011, so we're really trying to push and get ahead of this and make sure that these aren't just words on paper, but this is actually going to result in real improvement to WHOIS in the DNS space. I'm sorry that I talked so fast, but with that I'll give up the microphone.

MAGUY SERAD: Good morning everyone. Just a quick raise of hands… This is Maguy Serad, Contractual Compliance. How many of you have submitted WHOIS inaccuracy complaints? How many of you have received WHOIS inaccuracy complaints? That's it? Okay. Well, let's see what we've got to share with you this morning. We committed to the community last year that we're going to be doing some major overhaul in our compliance efforts to improve on the WHOIS inaccuracy processing, quality and turnaround time.

Basically this chart shows the efforts from September 2012. The reason we start in September 2012 is because that's when we really rolled out what we call a common process and a consistent process. To address contractual compliance complaints, focus first on WHOIS inaccuracy. With that process, this chart on the right shows the average turnaround

EN

time. The efforts focus started in September 2012 on reducing the process steps. In the past there were five. We reduced them to three.

We added more fact-based processing to where we're asking for very fact-based proof to be able to make a decision if this inaccuracy complaint has been addressed properly. The entire team has been cross-trained. We also put forth the effort to automate, in March, the whole processing effort and we wanted to maintain a turnaround service-level agreement for that. What does that mean?

Regardless of the volume of complaints, can you see the red line here shows the average turnaround time from when a complaint is received to when it's closed. In the past we've seen as much as 169 days of a turnaround time to be able to close a complaint. That's due to everything I was stating before – lack of consistent process, lack of transparency, lack of automation, lack of availability of data…

So there were a lot of factors. We wanted to measure the success of all the efforts that the team undertook. The best story to tell here is that if you look at the turnaround time of the WHOIS inaccuracy complaint, we are meeting the Agreement and the efforts that we put together for the process. The registrars have an obligation to respond to WHOIS inaccuracy within 15 days, and that usually consists of our first notice.

Then in the compliance process, for the second and third notice we require a five-day business day turnaround time response. I'm really pleased to inform you that all the compliance complaints related to WHOIS inaccuracy are within that turnaround time. One of the commitments we made with ATRT 2 and to the community also, is that

an effort to improve on the user experience and add additional features to the complaint processing.

We built additional FAQs. Those FAQs are related to each complaint type, and if you go on the ICANN.org website, on the compliance page where you submit complaints, you will see that the user experience now has FAQs and then a 'Submit a complaint' aspect. All this is available in the six UN languages. Additional features were added. In the past you were able to only submit one complaint at a time.

The first effort launched was to give you the ability to submit multiples, so as a user you don't have to re-enter your name every time, but you are given the ability to submit multiple complaints in one submission. Then we took it to the next level. We have what we call high-volume submitters. We reached out to those and conducted a pilot to test what we call a 'bulk WHOIS submission'. It's a way to upload a file to submit multiple WHOIS inaccuracies.

It can be provided, but the reporters will have to agree to terms of use and be provided with a user ID and login to submit and upload that file, because it's a different approach to it. We've improved the complaint submission forms and we've created and conducted several outreach sessions with our contracted parties, by region, but also at a registrar level. All these outreach sessions were based on volume trends and the issues and challenges we saw per region and per registrar.

The objective of the outreach session was to make sure that we build on the experience and knowledge to be able to respond and maintain the turnaround time and the quality of the complaints. Increased community reporting – we do generate monthly updates and they're

very brief because it's only about the activities for the month; a highlight. Those are provided in six UN languages.

In the monthly update we also provide you with some statistical information. We have increased our publishing, and in here you see some other areas that we publish. But the thing that was available today to the community that was not before, as Fadi announced two meetings ago, was the metrics. We provide you all kinds of data on MyICANN. Please take a look at it.

I know people say, "You collect a lot." Well, we do, because we have to be able to maintain our operations, manage our workload and manage the issues, but we also made a lot of that data available to you, and we've been receiving very good feedback from the community members who've used it. It's mostly used by Working Groups and the PDP efforts. Your feedback is always welcome in that are if you need something else.

Post-survey – contractual compliance is providing a very specific service to the community, so it's important to hear back from each and everyone who uses the compliance problem reporting tools, to know about their experience. At the closure of every compliant we ask not only the reporter but also the contracted party to provide us input on their experience; not only submitting but all the way through to closure of a compliant.

From a process, from a system, from people interaction and general experience. We take a serious look at the feedback provided and we do that as part of our continuous improvement. In summary, process system improvements, as we're committed to the community, as we just reported. Fact-based decisions, improving the experience, improving

the reporting, and having the ability to receive ongoing feedback from the community about the compliance experience.

I added a last bullet in here. If you were not familiar with the audit program, it closed last year at a 99% compliance rate, which is a good story to tell. [inaudible 00:54:38] program is underway now and we'll be publishing results on that later on in the year.

MARGIE MILAM: Thank you Maguy. Now we're going to move to Mary Wong from the Policy Team and she'll highlight some of the policy issues relating to WHOIS.

MARY WONG: Thanks Margie. Someone earlier this week said when the Policy Team shows up it's Clipart time, so here we are. My name is Mary Wong, I'm a member of the Policy Team here at ICANN. I'm here with some colleagues. There is a bunch of things going on at the gNSO that we don't have time to go into detail on. We have a lot more detail on the slides than I'm going to talk about so hopefully there will be some information there for you.

Some of you have been around when the gNSO commissioned some WHOIS studies. I'm pleased to inform you that they are all nearing completion and some have already been completed. We'd also like to talk to you today about three PDPs that either have just been done or have just been launched within the gNSO on this topic.

On the WHOIS studies, of the several studies that were commissioned by the gNSO a few years ago, two remain. They are completed and awaiting public comment. The one that's gone out for public comment – and thank you all for sending in your public comments – is the privacy and proxy abuse study that was performed by NPL in the UK, and ICANN staff is now in the midst of reviewing and analyzing the comments received with the NPL team; hopefully with a due date of releasing it back to the community early next month.

This connects back to something that Mike was talking about. I'll speak a little bit about this study when we talk about the PDP. The last study is the WHOIS misuse study being performed by CMU in the United States. They've just finished their initial findings and we plan to release the study for initial public comment shortly after this meeting closes.

I did want to show you the hypotheses that were tested by NPL as well as a very quick summary of the findings. This is a fairly complex study with a lot of detail, a lot of findings, so if you haven't looked at the report and you're interested in privacy and proxy issues, I do encourage you to read the report.

Essentially, you see the two hypotheses there and the question essentially was whether or not it's true that a significant percentage of domain names that are used to conduct illegal or harmful Internet activities, are registered via privacy or proxy services. NPL found this to be true. The associated hypotheses was that if these domain names are used to conduct illegal or harmful activities, are therefore registered via a privacy or proxy services, is that significantly greater than domain names used for entirely lawful Internet activities?

It was very interesting that the finding was only partly true. I'm going to run you through the next bunch of slides at a galloping speed. These are some of the other things that NPL investigated in looking at the two hypotheses. It's interesting that not only were domains registered for illegal activity reliant on privacy and proxy services, they also used multiple ways to hide contact details; including phone numbers and the like.

However, an above-average number of legitimate domains also do not provide accurate contact information, so it was also very hard to contact some of those registrants. I'd like to highlight the fact that for those domains that we're calling legitimate, those that were selected for this study were chosen to mirror the domains that were chosen to measure the illegal and harmful Internet activities – and you'll see that in a couple of slides.

On the left you'll see those types of domains that were studied. These were based on terms of reference developed in consultation with the community. In NPL's study it also highlights some of the scope of activities, some of which were suggested by the community that were not studied. The work was divided into a number of different work packages, and this slide summarizes the findings that I mentioned in the two hypotheses that were tested.

You may review these at your leisure. The final study is the WHOIS Misuse Study by CMU and I won't steal their thunder by revealing the entirety of their findings. Like I said, the report will be out shortly for public comment and as you know, the research team will also present a webinar that will summarize the methodology as well as the results.

Since they're the experts I would encourage you to listen to them at that time.

Essentially, you may recall that this was the study that we asked them to perform and the hypothesis is there at the bottom. We're talking about illegal or harmful Internet activities and domain name abuse in that respect. I do want to highlight that of all the terms and definitions used in the various WHOIS studies, they were used consistently so that the results can be helpful to the community and to ICANN in trying to figure out next steps.

This is particularly in review of not just the PDPs but all the other work being done across ICANN, including some of the efforts that we've heard from today. Just to give you a sneak peek, there was a descriptive study in the form of a survey of registrants, registries, registrars, law enforcement officials and experts in the field. The response rates were fairly low for reasons that they tried to explain in the study, but you see that these were the three types of misuse that were found and reported.

The other part of the study, design an experimental study that tried to corroborate what was found in the actual survey. So if you look at the draft report when it's out, they'll also go into details about why this might be the case, including whether the gTLD was significant. Again, as I said, the terms and the studies were consistent. This as well as the privacy and proxy abuse study looked at domains across the top five gTLDs.

They will also report on whether the price of the domain had any relevance, as well as the type of registrant, and the effect of any anti-harvesting techniques that might have been used by registries or

registrars to counter the misuse that was reported. Next steps. As I mentioned, we plan to analyze the public comments for the privacy and proxy study very soon, and we look forward to your public comments on the CMU study, also very shortly.

There are three PDPs that are going on in the gNSO right now. The first I'll speak on is one that was recently kicked off, last month, by the gNSO. This had to do with the issues that Mike spoke about in his presentation, regarding for example, the current temporary specification in the 2013 RAA. Essentially, this PDP is going to be looking at the issues surrounding accreditation of privacy and proxy service providers.

As many of you know, ICANN has already committed to creating that accreditation, so the hope here is that the gNSO's PDP will lead to some concrete recommendations, suggestions, possible best practices, for what such a program might look like. Like I said, this group was only very recently chartered.

We just came from the first informal meeting of the Working Group and we also discussed some of the ongoing work in the rest of the community, in particular some of the issues, topics and research that they highlight, most particularly in the recent status report that they just released. The reason I'm mentioning the PDP now is we would like to encourage those who are interested in the topic, who are not yet members of the Working Group, to join us in our work.

The next PDP is one that's nearing completion and that is on Thick WHOIS. As I think everyone knows, WHOIS requirements are specified through registry and registrar agreements, and there are two models; one is Thin, one is Thick. The difference really is in terms of whether the

extent to which registries maintain, store and provide these sets of data that are associated with WHOIS and the domain name registration.

This PDP was chartered to examine, in part, the consequences of requiring Thick WHOIS across all the registries. There were a number of public comments that were received in response to the report that came out of this PDP and the Working Group. The Working Group achieved full consensus on its recommendations, and this essentially… Here you see what the report recommends, published last month.

Since this report and its recommendations achieved full consensus in the Working Group, it was also approved and adopted by the gNSO Council, so essentially when we say nearing completion, it means that we're now moving into the implementation phase for this particular effort, and there will be an Implementation Review Team.

In fact, there is an Implementation Review Team that's been created, which will allow the community to work together with ICANN staff to oversee the transition and implementation of those consensus recommendations requiring or recommending Thick WHOIS across all registries, into the implementation phase.

There is a public comment period on the implementation phase, so if you have any comments or suggestions, please let us know. The public comment period is still open, through I believe early December. For further information this is a link to the public comment forum, as well as to the report if you haven't already seen it. At this point, for the last PDP that we thought relevant to the discussion today, I'm going to hand off to my colleague, Steve Sheng.

STEVE SHENG: Thank you Mary.   Just some background – PDP translation and transliteration of contact information – the Council adopted their Charter yesterday at this public meeting and a call for Working Group volunteers will be published shortly. The Working Group is dealing with two issues: should local contact information be translated into a single language, for example English, or should it be transliterated into a single script.

The one-sentence policy question is how does the community balance the capacity of registrants who may not have the capacity to enter WHOIS information in for example English, and balance that with the potential global consumption of that data? The second question is, who should bear the burden to either translate or transliterate the contact information?

This issue was first caught out in 2006 by an Internet Architectural Board document, and has been worked through the ICANN community; first through the [RAD? 01:07:37] Working Group and now through this PDP. To help…   As part of the effort to provide objective data, there is a commercial feasibility study commissioned on the transliteration and translation of contact information; to look at how much does it cost, some of the accuracy considerations for that. That's a very high-level overview of what this Working Group is doing.

As I mentioned earlier, the Working Group will assume its form and there will be a call for community volunteers.   Staff will also be conducting a webinar to provide background to the community on this issue, and I invite community members to really participate in this PDP.

These are the places of where to find the relevant information. Next I'm going to hand over to Jim Galvin to speak about the WHOIS Review Team.

JIM GALVIN: Thank you Steve. I'm not actually going to speak about the WHOIS Review Team, but this is an EWG that was motivated by two recommendations that appeared in the WHOIS Review Team set of recommendations, and they're listed up here on the slide but you can also go back. I encourage you to look at the complete text of the WHOIS Review Team report to see the full context here.

This group is relatively new. We've only had three or four meetings and we've only completed one particular action, which was to agree in the Working Group as to what our Charter is. We had a call for volunteers that was published in July and the group was formed in September based on a selection of eight volunteers. We had people from around the world as well as from the various parts of the system; registry, registrar, technical expertize and some policy folks, who are part of this EWG.

The group has two deliverables that we're looking at. One is to focus on the requirements for internationalized registration data. This particular deliverable comes directly out of the Internationalized Registration Data Working Group that closed a couple of years ago. The second part is to consider a data model for that internationalized registration data. These again were recommendations that came from the Internationalized Registration Data Working Group as well as being related and motivated by WHOIS Review Team recommendations.

Of course the staff is putting together this Working Group as part of its action plan on what to do about updating WHOIS. We now have a location where you can go for the community to look and track the work of this group. The Charter document is up there and the list of members. As I said, we're just getting started. We've only had a few meetings and our focus so far has been on carefully scoping out our Charter according to the request from ICANN, as part of its request for volunteers.

One last thing that I would like to say to get ahead here, there are obviously a number of active Working Groups in this space of WHOIS updates. There are a number of moving parts so I thought I'd take a moment to talk about what I think my relationship is with some of these other groups. In particular we have the WEIRDS update, which is not an ICANN group, per se, but it is an active Working Group.

I would expect that the output of this group would help to inform that activity, since part of the reason for having the WEIRDS group and its work is to look at how to support internationalization of registration data. Once we have our requirements set, I hope that would inform that activity. We also have the Translation and Transliteration Working Group, which has just started. I expect that work to be complementary to ours and to work closely with them.

We will have a set of requirements. We're very interested in the results of that work and the process by which they get to their decision of whether or not there will be active translation and transliteration, and where that will be. Our focus will be on the presence of data that would

need to be translated or transliterated. We want to make sure that our requirements, anything we want to go forward with, are not conflicting.

In the end I'd expect our work product to potentially inform a future PDP process. We will be simply specifying a set of requirements that we'll advise to the community, and that would of course go for public comment, and we'll see at the end of that what our specific suggestions are for recommendations, but I'd think the next step could potentially be suggesting a PDP process that moves that forward. Thank you.

MARGIE MIRAM: That concludes the presentations. Now we'll have questions on any of the presentations. Please stand at the mic, give your name and we'll address the question to the appropriate presenter.

STEVE METALITZ: Hi, Steve Metalitz with the Intellectual Property Constituency. Thank you for this presentation.   It was very comprehensive so I really appreciate that. It helped to put all the moving parts together a little better for me. I have two questions and I think they're both for Mike. One on the cross-field validation issue – which I understand from the transition specification that that requirement doesn't come into force until ICANN and a working group of registrars agree – what is your timeframe for bring that provision of the agreement into force?

That's the first question. The second question is there are several places in the WHOIS accuracy specification in which registrars are told that if they need to verify or re-verify some data, and if they don't get a response from the registrant they either have to suspend the registration

or manually verify the data. I wonder if you could tell me what is the definition of 'manual verification' in that context?

MIKE ZUPKE: Great, thanks Steve. The first question was about the timing for the WHOIS address validation. We, meaning ICANN in this registrar Working Group, we worked pretty hard on this before we even had the RAA approved by the Board. We spent a lot of time on calls and exploring exactly what was involved and talking to vendors. So quite frankly, when we got the Agreement approved, we put all of our staff into getting registrars on the Agreement.

That got deferred a little bit so the plan right now is to pick it back up after this meeting, when the staff starts a new trimester. My own personal goal is to get that back on track and come up with a project plan that we and the registrars can all agree to and say, "This is a reasonable timeframe that we think we can conclude this." So I don't have an end date but I can assure you that it's on the radar for the coming trimester; to get it back in motion.

To your second question, I think there's probably not a clear answer to that. I think this is an area where we said we were doing something sort of new and we'd learn as we go. Maguey or people on her Team may have some ideas on this, but I think the recognition is that automated processes are not going to be 100% reliable, and so the idea is we don't necessarily want registrars to flush customers away just because they have difficult data to automatically validate or verify.

So I think there's generally an understanding that there's got to be some reasonable business flexibility in that, but I don't have a more specific answer. I don't know, Maguy, if there's something that you wanted to… I think one of the things we're doing is we're rolling out the new RAA – I mentioned we've got the ticketing queue for registrars – is we're posting the FAQs, we're doing these webinars, and we're also trying to post advisories on the new topics.

You may have seen we posted an advisory to registrars back at the end of July about what some of the transition timing meant for them. So this could end up being a subject of an advisory if it seems like there's lack of clarity about what that means.

STEVE METALITZ: Thank you. You have a ticketing system for questions from registrars. Do you have a ticketing system for questions from non-registrars?

MIKE ZUPKE: My personal email address! In honesty, I'm always happy to take those sorts of things. I think lots of us on staff get them. The idea for this particular ticketing system was we wanted to give consistent answers to registrars; that they'd all be at a level playing field and nobody would have an advantage because we gave them incorrect information, for example. But I'm always happy to… I think I may have put my email address on one of my slides, but… Thanks.

AUDIENCE MEMBER: [inaudible 01:18:12], Chair of the registrars, I'm speaking in that capacity. Things like some of the terminology, Steve, just to address 'manual verification', if it has to be vague; you have to allow us to verify using whatever means we deem are suitable – it's a contract so if you put in something too specific then it would preclude us from doing other things. I don't think that would be reasonable and that wouldn't help you either. I think the fact that it might be a little bit vaguely worded is better for both us and for you.

I think some of the new provisions within the RAA are things that are going to cause headaches for both ourselves, ICANN and the rest of the community, as people try to understand exactly what the terminology means. That's going to be an interesting process. Just for the record, with respect to WHOIS and everything else, it's interesting to see that on one of those slides you're talking about doing a legal review to see whether Thick versus Thin WHOIS has potential implications for privacy law.

Yet, you're still trying to force registrars to sign a 2013 RAA that puts them in breach of the privacy law. I thought that was rather amusing but that's just a personal comment. Thanks.

MARGIE MILAM: Any other questions or comments? Do we have anything from remote, [Charla? 01:19:49]? Well, thank you very much for your time and for listening to this update. We plan to provide updates going forward. Thank you very much. [applause] One of our mandates is to keep the community informed on our progress, and so over time we'll have

ICANN 48 · 17-21 Nov 2013

additional updates and presentations at future ICANN meetings. Thank you very much.

[END OF TRANSCRIPT]