SAN JUAN – Identifier Technology Health Indicators
Tuesday, March 13, 2018 – 17:00 to 18:30 AST
ICANN61 | San Juan, Puerto Rico

CATHY PETERSEN: Good afternoon, everyone. We will be starting the Identifier Technology Health Indicators session in a few minutes. We'll just give it a couple more minutes. Thank you.

ALAIN DURAND: Good afternoon. This is the ITHI session for Identifier Technology Health Indicators. This is a project that has been started for a while, and today we are going to show some interesting numbers [inaudible]. Number that will be interesting to you. They were interesting to me.

In this session, we will have three presenters. The first one will be Paul Wilson the current chair of the NRO. He's going to give us an update on what the NRO has been doing in this area.

The second one and the third one will focus on the main part of the project, things that are managed by ICANN. The second presentation will be made by Christian Huitema on the current metrics and the data that we are finding on the current metrics.

The last presentation will be made by Geoff Huston on proposed set of [some] new metrics.

So without much further ado, I will leave the floor to Paul who is going to talk about the activities in the numbers space. Paul?

PAUL WILSON: Thanks, Alain, and hi, everyone. The regional Internet registries all of us run five, all of us run a WHOIS registry using a fairly familiar WHOIS service. So there are five different registries which are run by the five RIRs. They're quite tightly coordinated amongst each other. Technically they're capable of inconsistency and, of course, 0f errors and incompleteness and so on. So the RIRs work together under the banner of the NRO to make sure that those registries are making sense relative to each other and that they're fulfilling their purpose in terms of the records that they store.

We've done that for a very long time, but I think things are changing somewhat in recent years in there being a much higher interest from a much wider group of relying parties in the databases in the correctness and effectiveness of the databases. Also, the pace of updates is increasing quite a bit as well. So before we hit the current shortage of IPv4 addresses, allocations were fairly static. They were made to parties who kept those allocations and used them. But these days, we have a lot of

transfers happening. So within the regions and between regional Internet registries, there are a lot of transfers going on, which obviously require updates to the database. So that's another reason why our focus on the correctness and completeness and so on is increasing.

As I say, we have been concerned ever since the inception of the RIRs really that the registries do their job. We haven't spoken about health as such, but the idea of identifier health is something new that I guess has come with ICANN's project. But that said, we have maintained this focus.

The other aspect to this, of course, is that we have membership relationships with the network operators who are the first recipient of address blocks and ASNs. So they are obliged under their relationships with each of the RIRs to keep their records up-to-date and current. There are policy issues there in terms of what exactly are the expectations and the penalties, so to speak, for not complying with those policies.

These things are being handled at a regional level generally. So the five RIRs do have independent policy processes and memberships and will at different times have discussions about WHOIS related policies. Those discussions, as I mentioned, are also increasing in frequency and intensity these days. I'd say that around the five regions, I think it's fair to say that we're all

tightening up in different ways but in generally the same direction toward clearer and more tightly implemented policies.

The ITHI project really came as no surprise from ICANN. It's obviously a shared interest of all of us who perform registry roles, and so we decided to plug into the ITHI initiative of ICANN's. I think it was actually quite useful because, although we've worked together quite closely, we haven't actually settled on a set of consistent metrics which we have now moved toward through the ITHI project.

We actually spent a period last year, went through our registration services coordination group that's a group of staff across the five RIRs who work in the registration services areas. That group did some work on a draft set of metrics for what we would refer to identifier health in the numbers space. They also launched a public consultation on a draft paper.

That happened toward the end of last year, and it gave a chance for our communities to feed back into that process. We actually got very little feedback, so we now have a document that's close to ready for final approval and publication. It basically documents the identifier health in the numbers space in terms of WHOIS records. What we've arrived at is the three Cs: comprehensive, current, and correct data being our aim.

But those are broken down into five specific measures which are measurable metrics of our database which relate to the completeness of the database, uniqueness, the matching of our database with other external official records, the effectiveness of the data concerned in actually reaching people who are documented in the registry, and the up-to-datedness of the data as well. That document also identifies the various risks associated with not reaching our targets in those measures, and it analyzes the causes that would be associated with that kind of failure.

So that document is going to be released shortly. What we don't have yet, obviously, since the metrics themselves are still in draft form, we don't yet have what I think Alain is going to present shortly, which is actual data about our compliance. But obviously the point of having these metrics is to be able to measure stuff to set some targets with which we hope to comply and to track the degree of compliance over a period of time.

So I think if you're interested in the numbers space, then watch this space and we'll be able to report in due course on the health of Internet identifiers in the numbers space across the five RIRs. I think that's all. Thanks, Alain.

ALAIN DURAND: Thank you, Paul. I would like to take this opportunity to thank you and the other members of the numbers community for working with us on this project. I think it's an interesting collaboration, and we are learning a lot going through this process.

PAUL WILSON: Yeah, I agree. Same to you, Alain. Thanks.

ALAIN DURAND: Now we are following the same process like looking at the [problem] space and defining the metrics and then getting to the actual measurement so it is actually normal. But you don't have the numbers yet, and we will be looking forward in the future to see the first batch of numbers whenever those will be ready.

Now we're going to shift gear and move into the namespace. Christian is going to make a presentation on where we are.

CHRISTIAN HUITEMA: Good afternoon. I'm Christian Huitema. I have been working on the measurement of the DNS data and [status] for the last year and a half, about, after doing a study to see what could be done with the DNS working with Alain on doing actual metrics.

As Paul said, there have to be some principles when you set up for those metrics. The principles that we have adopted are pretty much on this slide. It's first that we really wanted this metric operation to be technical. We don't want to be involved in policies. The purpose of the metrics is to describe the state of the system. It is not to basically make judgment one way or another.

We have been looking at defining areas that we want to track that are potentially problematic, defining the metrics in these areas, and defining ways to measure them.

Another principle is that we don't want to take snapshots. What we want to do is to have a continuous system that operates for a long time and basically gives the metrics what we are targeting every month publish a new value and, of course, publish the value of the past months as well so that we can estimate trends. Because we believe that trends are almost as important as the actual value.

That's the reason why in doing that we are investing a lot in automation. Basically, we are setting up probes at various places, and we are doing constant feedback and automation. So the website that publishes data is automated, so the metrics are automatically produced every month, etc.

In the slides that we are presenting, we are going to give you the measurements. And for the same reason that we don't want to be involved in policies, we want to give the measurements as they are. Each time you see a number, you say, "Oh, Widget X is now at 29%. Why is that?" Well, our generic response is, "We do not know why that is." I mean, we have guesses, but your guesses are pretty much as good as ours. So we don't want to put those guesses in the metrics publication. The metrics are straight measurements.

Another principle is that we are very careful to not have privacy issues. So all the data that we are publishing are statistical in nature. All our tools are open source, and all our results are published so they can be analyzed.

We had a couple of presentations in previous sessions. We had the presentation of the ITHI metrics already in Abu Dhabi, for example. They are in seven categories for us. One, we are looking at the accuracy of the WHOIS data. We are looking then at the behavior of the root servers and the level of abuse that they are taking to some degree. Excuse me, the domain name abuse, abuse of the domain name system. We are looking at the DNS root traffic.

For all of these metrics that are listed here, we have sources of data. For example, for WHOIS, we are working for the ICANN

Compliance department. For domain name abuse, we are working for the DAAR project. For the measurement of root traffic, recursive servers, IANA registries for DNS parameters, and DNSSEC deployment, we are working with scans of root traffic or with scans of recursive resolver traffic. And we are collaborating with recursive resolvers to [probe] effectively that get us those statistics.

Timeline, we have been working in the past year on the definition of the metrics. What we have now is a presentation of the first data. In the last two months, we have set up the initial captures, and we have been able to get data for M1, M2, M3, and M7. Geoff Huston will present the data for M5 in the next talk. We also have been able with early collaboration to get an initial set of data for the metrics M4 and M6, which are about client use of the DNS.

So we integrate M5 as it is developed. We are going to build up the pipeline and get more probes so that we have data that is richer from the metrics M4 and M6. We are going to enrich that and publish that on the ITHI ICANN website.

First metrics, M1. M1 is tracking the accuracy of WHOIS data. We are doing that by using a proxy for accuracy, which is the number of complaints. We don't take the [real] number of complaints. We take the number of complaints that have been

validated by ICANN's Compliance department. Right now, that number stands at a little bit less than 6 per million. That's our first data, so we don't have a trend yet. But we are going to track that trend over time.

With all this data, what we see is that the average doesn't tell the story. If I tell you there are 6 complaints per million domain [numbers] registered on average, well, that's just an average. We have plotted here a curve which is the [inaudible] frequency of complaints. Basically, the total of complaints in the Y axis and then the X axis the number of registrars ranked from the one with the most complaints to the one with the least.

What we see there is that the distribution is not [inaudible]. If it was every registrar has as many complaints, you would see a straight line on the diagonal. That's not what you see. What you see is the line is very curved, very tilted toward the Y axis. In fact, it takes six registrars to account for at least 50% of the complaints. I mean, it's not an even number, six registrars account for a bit more than 50%. It takes 44 registrars to account for 90% of complaints. That's on a total of almost 2,000 registrars. So there's a very [skewed] distribution there.

As I said, these are [inaudible] number. It's not a judgment or a reasoning of the cause. But that's what we observe.

CATHY PETERSEN: Excuse me, Christian. We have a question online.

CHRISTIAN HUITEMA: Yes?

CATHY PETERSEN: From Kathy Kleiman, "How do you know that WHOIS complaints are valid? We understand that some are done for harassment purposes."

ALAIN DURAND: I will answer this question. We have been working closely with the ICANN Compliance department. We are not looking at all the complaints. We are only looking at complaints that are related to the accuracy of the data. There are many other types of complaints that we are not taking into consideration.

The ICANN Compliance department has a process where they look at those complaints and evaluate them. If they think that there is enough ground for those, then they send what they call a first notice. If there is no answer, then they go to a second notice, and they go to a third notice, and then potentially all the way to a [breach]. So that's a process which is very well defined, that is well documented in the ICANN Compliance department.

So to answer this question again, we are only looking at the complaints that are related to that accuracy of the WHOIS database of a registration and complaints that have been validated, that have been through the first notice stage.

CHRISTIAN HUITEMA: Thank you, Alain. So that's the metric M1 about the WHOIS accuracy. The series of metrics M2 are about the abuse of the domain name system, and we are working with the DAAR project for that. They are tracking four types of abuse: the number of websites used by phishing domains, the number used by malware domains, the number of botnet command and controls, and the number of spam domains. The metric is defined as the number of abused domains for 10,000 domain names.

We see there the global averages, which are basically on the order of 4 or 3 for first three types of abuse and of a much larger value for the spam domains because spam is a very widely distributed activity.

Now we also in the same way that we sorted for M1, we also see that those averages don't tell the story. If I look at the distribution per TLDs, we see that for example when phishing is concerned a single gTLD accounts for over 50% of all phishing domains. And it takes only 11 gTLDs to account for all the

phishing domains. We see the same kind of skewed distribution for the other domains. So that's clearly an indication of the structure of the problem.

We have been trying to do the same measurement for registrars, but we don't want to spend too much time on the registrar data there because our registrar data has to be evaluated with the WHOIS process and it is subject to all the restrictions of using WHOIS data as in limitation in time and all that. So we only publish seriously when we get data that we can actually verify that we can trust, and today it's a bit preliminary.

But we intend to present this skewing of the data in some kind of a table like this one that says, okay, how many gTLDs does it take to account for at least 50% of the phishing domains, of the malware domains, etc. And then how many does it take to account for at least 90% of those variations. As I said, we're in the business of measuring stuff. We make no interpretation, and we make no reasoning about why it is so.

The M1 and M2 data are produced by the Compliance department of ICANN and by the DAAR project, and they are about the quality of the data. The M3 and M4 data that we will see later are about actual DNS traffic, what we see there. M3 is about root traffic. We measure the root traffic by instrumenting the L-root. We are basically doing about one sampling per day

per L-root server. Those samplings are taken at random times, so they account for all the time variations when we aggregate them statistically. Then we are getting all those samplings and summarizing them every month and getting those metrices.

What you see there is the first metric: how much of the root queries get a "no such domain" response? And it's pretty large. That's effectively almost two-thirds of the root traffic, queries that have no particular value. Then on the remaining queries, we look at how many of those queries could have been cached by the resolver. Again, we see that it's a fair share, almost 30%. The queries that we don't know whether they could have been cached, they probably could not be, is the in the order of 6-6.5%. We track that every month. You see here the current value and the average and the pie chart that shows how they split the domains.

For the "no such domain" queries, which is this big part of the pie on this circle, we tried to slice the pie into components. What causes that? We have found we are looking at four components: the reserved names, the names that have been reserved by the IETF like .local for example, and there are five or six of those, which account for about 3.4% of the traffic; the frequently leaked strings like for example .home which account for 9.3% of the traffic; and the frequent patterns, we see a pattern in the data. They are not frequent strings. Each name appears only a

very small fraction of the time. There are many, many different names, but the follow patterns and we try to account for those patterns. And then everything else. There are about 10% that we cannot directly explain by either of those processes.

For the special use names defined in RFC 6761, what we see is that the bulk of the usage is with the .local domain. It's about 2.77% of the traffic on the root today. Other reserved domains are present but are present in much smaller numbers: .localhost is present quite a bit, .invalid is present quite a bit, and then the other domains are really traces.

On the frequently leaked strings, what we are doing here is that we are getting the strings that are most frequent at the root and in the current variation in the current implementation we are only looking at strings that happen at least .01% of the time.

In this particular slide, I only give that strings that happen at least .02% of the time because the lower the number, the less sure you are of the results. And also because it would make the PowerPoint very hard to read.

Again, we see that there is one name that dominates that, which is .home which accounts for 3.5% of those requests that the root sees. Then there is a series of other names. The take home thing there is that we can absolutely measure the leakage of those names at the roots, and we can track that month by month, and

we know which names are being used and which names are recurrent. We may see that it changes a bit from month to month. Some names are going to disappear, but we can see that there's a core of well-used names that appear all the time.

I told you that a number of the names that we see in this root traffic are not special use domains and don't correspond to frequently used strings. They are just random names. In fact, if you see here in this distribution, we did a distribution of those names by length. We see the bulk of those names are in the 7 to 15 character long. The longer names we did not plot because there are very, very few of them.

Many of those names with the length between 7 and 15 when I look at them by doing a random sampling, they look like things that have been randomly generated by computers. They are not all like that. It's actually very hard to distinguish what is randomly generated by a computer and what is just some kind of numbering plan somewhere in a Wi-Fi network, for example. But it's something that we want to track and we want to go further and analyze that further.

That's the traffic at the root. Now when we did that first study last year, we did some experiments and we came very quickly to the conclusion that the root was not necessarily representative of all the traffic by the users. If you understand the DNS

architecture, you know that what is seen as the root has already been filtered by DNS resolvers. In fact, if the DNS resolvers applied all the modern technology defined by the IETF, you would see extremely little traffic at the root. They would cache the good results. They would cache the [unsatisfactory] results. So we would see none of that. So a lot of the traffic at the root corresponds to anomalous behaviors.

If we want to look at what users are actually doing, we want to be close to the clients. That's why we have been working with recursive resolvers to put probes at recursive resolvers and try to look at what's happening there. How much of the queries issued by clients are going to registered TLDs rather than all these strings that we see at the root? How much are going to these IETF reserved names? How much are going to the frequently used strings that we see there and what else?

Now you remember that when we are looking at root traffic, we see these queries to nonexistent TLDs represent almost two-thirds of the traffic. Here in the one probe that we have – and I must qualify our data that we have only one point of measurement today. We are ramping to get more. In this one point of measurement, these nonexistent TLDs represent just 1% of the traffic, much fewer.

The trends are also different. In the reserve names where we see a small amount of traffic with .localhost, .local, and almost nothing for the other names.

In the frequently used strings, that was a bit of a surprise for us. It's actually dominated by local names, like host names that people try to resolve and they don't do their queries correctly and they end up sending the [query] putting the host name as a single [token] name that could be mistaken for a top-level domain. We don't publish the value of those names because there are privacy issues. They are typically names in the local infrastructure of the people that are [providing] the probes, so we put all of them in just a global category of "local host names."

If we go beyond that, what we see is very little traffic for these kinds of names that we see at the root. We see some traffic for the big names like .home, but we see traffic for names like .dns, .internal, or .unifi, which in that case represents the Wi-Fi network that they use. So that was one of the lessons for us. At that point, we want to have many more probes before we can make definitive statements but we see that there is a difference between the traffic at the clients and the traffic at the root.

You wanted to intervene?

ALAIN DURAND: I would like to add a little point to what Christian just said. So far, we have been working with a number of small organizations and we already have two organizations that have agreed to participate and are already contributing data. I would like to recognize them here.

One of them is the University of Cape Coast in Ghana, and another one is the University of La Plata in Argentina. We also are now working with a third organization Nawala which is [a more or less] service provider in Indonesia. Last night, we were up very late trying to help them to install the tools to do all those measurements.

We are reaching out to other potential partners, and our goal is to get some more participants to this. If we could get maybe five, six, up to ten maybe by the end of the year, we'll be quite happy. We would like to get different types of players, some that will be academic, some may be more industrial, some may be service providers, some may be small or large or very large.

But we are starting with a core [inaudible] approach. We started small. That has enabled us to understand how things really worked to finetune the tools that we have. Now we are developing a process to make this more automatic. We can go to larger players, and hopefully even at some point to much bigger players.

So I wanted to thank Christian for [writing] the tool and helping everybody to deploy it.

CHRISTIAN HUITEMA: Thank you. Well, Alain actually spent a lot of time in deploying them too. The whole point of this worldwide infrastructure is that you spend a lot of time in phone calls or computer chats in the middle of the night. But that's part of the territory, I would say.

So M3 and M4 are analyses of two parts of the traffic. What type of DNS traffic do we see at the root and at client side? With the M4 data, what we wanted to also see was how good and how useful are all these IANA registries that we are doing for the IETF? We cannot track all the IANA registries because we have only DNS [related] data. But what we did was for the DNS [related] tables, look at parameters that are part of registries. For example, the [r] types or the [r code] classes, but also parameters used by DNSSEC or parameters used by DANE.

For those parameters, we wanted to answer two questions. One is, do people actually use the data that are registered? Basically what we did, we said, "Okay, if a table defines ten values, how many of those values do we actually see at least once in our data set?" For some tables, the answer is zero. There are a few tables like that.

For the classic tables like the DNS classes or the algorithm numbers, the answer is between 20% and 70%. Some values are seldom used. For example, in the security algorithms, some security algorithms are obsolete and people don't use them anymore. But we can see that. That gives us an idea and a confidence that what the IANA is doing is useful.

The other thing that we wanted to see is whether people were bypassing the IANA registration and directly making up their own values. In that set, we only see that for the DNS option codes, the EDNS0 DNS option code, while there's some usage of experimental values that we see in the wild. So globally that's what it is.

Now I would like to make a note there about the TLSA certificate usage and generated DANE certificate. In my data, I don't see them. So I had a long conversation with Victor Dukhovny about that. He told me that was normal because most of the DANE usage is between a mail server and authoritative servers. The mail server will be querying the authoritative server directly. So that traffic will not be caught at our probe points. I'm working with him to get a direct feed of the traffic that he has on his DANE measurements so that we can actually evaluate properly the usage of the DANE tables.

This basically is the way we can use those measurements to track IANA. We don't track just one table. I gave the data for four or five tables in the previous slide. Here that is the whole list that we are doing there, and we might add eventually more tables to the list when we figure out how to parse the data and extract them.

The final metric, M7, is about DNSSEC deployment. We started this evaluation of DNSSEC deployment by parsing the root zone to see how many TLDs were providing a DNS key. That number is quite stable at something like 90%. But we hope that it will change over time and reach 100%, but it changes very slowly.

Now analyzing the M4 data, we realized that we were seeing a large part of traffic that was actually DNS security traffic. We see that because we can notice when a client is using DNS security they place DO bit in the queries that [inaudible] in the response. So we can measure the fraction of queries that have that bit and say, "Hey, if we can find the client doing that, we know that so many clients are using DNSSEC."

So we can add to this data what we want to do in M7.2, which is the percentage of DNSSEC queries from clients that are using DNSSEC. If we are really ambitious, we certain also see the percentage of queries from recursive resolvers that are using DNSSEC and, interestingly, the percentage of responses from

authoritative servers that are providing DNSSEC answers. I think that by doing that, we'll get a handle on actual DNSSEC usage and be able to answer the question, how much of DNSSEC is used today? I think that's something that would be interesting for the community.

So I've been going through six of our seven metrics. Geoff Huston will present Metric 5 after me. M7, as I say, is very stable, so this kind of graph doesn't tell us much now.

I would like to thank you for your attention and answer any questions that you have now if you have questions.

RUBENS KUHL: Rubens Kuhl, .br. I'd just like to comment that the recursive server, recursive DNS, recursive metrics are being based out of three recursive servers. And we currently have 50,000 [inaudible] systems on the Internet, so publishing those results until we get at least 5,000 DNS recursive servers probably not what we should do since it has no statistical relevance whatsoever. It's like putting one [inaudible] on the microscope and deducing all fabric in the world based on that.

[So it amuses me] that such a metric would be published by ICANN, especially in an area where ICANN has no direct data on different from the authoritative root data when it runs one of the

most comprehensive root server systems because it's [inaudible] instance. So a root has a very good statistical significance among the root queries. But as for recursive queries, we shouldn't publish it at all until it crosses a really good threshold of statistical relevance.

CHRISTIAN HUITEMA: That's indeed a very good point. We have been using all kinds of caveats in this talk to explain that we have only one point of measurement now and explain the [inaudible] we want to do. It's clear that we want more than point. I don't know that we need 5,000. I would be happy to have 5,000, but I don't know that we need 5,000.

What I plan to do is to compare the data from the many sites as they sign on to see how they differ and how they are common. The idea being that we know there are differences. There are differences in time, like in the morning and the evening are not the same. In the weekend and in the workday are not the same. We know that there are differences in geography. People don't ask for quite the same traffic in China and in America. We know that there are differences in type of occupation. People don't ask the same query in an academic and government or in an enterprise or in a private network or in a mobile network. So clearly we want to have representation of all of those.

The statistical significance is something that we will address. That's definitely something we want to do. But we have to start somewhere, so we are effectively collecting the data. And we will be comparing the sources so that we can actually answer your question.

RUBENS KUHL: Yes, but I would like to respond to that. While those caveats are described, they usually go as very small letters. So anyone that reads what's in the report will repeat that and publish that in press and social media and not reproduce the caveats that this data is actually meaningless. So actually publishing that is a disservice to the community. That's my point on that.

One comment I had on another matter is that it was mentioned that some 0f the registrar queries were affected by limitations in WHOIS and so forth. There is data that I can collect from all [Thick] registries which is the BRDA which is the Bulk Thin Registration Data. That Thin registration data already contains which registrar is associated that domain to. So there is no need to do WHOIS query. There is already data inside ICANN that provides that information with 100% precision, so you might want to look into that as well.

CHRISTIAN HUITEMA: That's an interesting point. Actually, it's a good point. I would like to address it by first saying the ITHI project is a customer of the DAAR project. We are getting data from DAAR so whatever decision DAAR has been making, we are inheriting. So first I would like to suggest you redirect your question to the people running DAAR.

The second [and I'll somewhat] try to channel them. From my understanding is they wanted the study to be replicable, meaning somebody from the outside not ICANN could actually replicate the exact same study, open methodology and data that are accessible. The data that you mentioned may or may not be accessible from the outside, and that would put ICANN in a unique position to be the only one being able to do this study. Their choice was to not go that direction. They may change their approach at some point, and maybe John Crain can say a word about that, but so far that is the direction that we have taken. And as a customer to them, we are inheriting this decision.

So the question is, why do we have to rely on WHOIS to do attribution to a registrar as opposed to using data that's internal to ICANN?

JOHN CRAIN: If we have all of that data available internally, I've not found it internally. That would be awesome. But one of the things we

were trying to do is make it replicable by other people, which means using external sources. The only thing we need out of WHOIS is the registrar ID. We were actually talking earlier about where there may be sources internally, so we may actually have to swap to that because I think the WHOIS just may not be practical.

RUBENS KUHL: The source is actually called BRDA, so you might want to look into that or hack those servers and get the data out of them. But even if you use that data, that still makes things reproduceable but just makes it more complicated for other people to actually use WHOIS queries. But they can reproduce using WHOIS queries because it's the same information. So it's not privileged information in any way.

JOHN CRAIN: Yeah, understood. When we started the project, we were very much everything should be exactly as people outside would do it. And you're right about the reproducibility, so we are reconsidering.

RUBENS KUHL: Okay.

ALAIN DURAND: All right. Are there any questions on the chat room, Cathy?

CATHY PETERSEN: No.

ALAIN DURAND: No questions? Okay. So then thank you, Christian.

CHRISTIAN HUITEMA: Thank you.

ALAIN DURAND: Thank you very much for showing the numbers for the very first time here. Now I would like to invite Geoff Huston. Is Geoff in the room?

CATHY PETERSEN: Over here is a question.

ALAIN DURAND: Oh, we have a question.

SEBASTIEN BACHOLLET: As Geoff is not in the room, I just wanted to ask one question and from somebody with no technical background. Is there any link between what you are doing and some of the questions about they key rollover and the data they need to understand what is happening? Really sorry, it's [inaudible] question from [inaudible]. Thank you.

CHRISTIAN HUITEMA: As of today, the answer is no. We have no connection. This is not a metric that we considered initially. Now as in the future there might be more rollovers and they may be more or less frequent, it might be something that we would like to track. So today we talked about seven metrics. We think that we understand them well enough to be able to measure them. We are thinking now about the second phase where we'll add more metrics, that we'll look at other types of problems. And that one may be one of the areas that we need to look at and add to what we are doing now. [responds in French also]

SEBASTIEN BACHOLLET: I get your point. I just want to be sure that my question was well explained, and sorry for that. It seems today that we are missing data to be sure that it's the right time to do the key rollover. It's not the fact that when we will do a key rollover, it's each year you will be able to gather data [as a question] that if with your

project data there are data who can be used by the people who need to decide when to do the key rollover.

CHRISTIAN HUITEMA: As of today, we do not have data that will help them.

SEBASTIEN BACHOLLET: Thank you.

PAT KANE: Hello. Pat Kane with Verisign. Just a follow up to Sebastien's question. Earlier today, the CTO for ICANN correlated a decrease in DNSSEC queries to the KSK rollover push from last fall until next October. So I think it's important that we understand from the usage of DNSSEC how that relates to it to inform that decision because a lot of the data in terms of people with resolvers that don't have both key pairs is getting worse than it was late last year. So it would be very good to get that information to David sooner rather than later. Thank you.

ALAIN DURAND: Thank you. It's a very good point. As we have seen, Christian was talking about a new metric M7.2 that will actually track the number of queries that have the DO bit set. That may help in that direction with other metrics that we are trying to design in

that particular space. So maybe we can have an offline conversation if you have specific ideas on what we should track.

Is Geoff back in the room? Okay, so I apologize. We lost one of our speakers. I can just talk briefly about what we are planning to do.

Metric M5 initially was one of the metrics looking also at the resolvers but more from a client perspective. We have asked Geoff to look at this, and Geoff has a system of measurement that's based on Google Ads that's well-known and we have been using it in other contexts. We have asked him to use the system to explore what can be done from the perspective of a client, the stub resolvers.

One of the things that we would like to look at is are resolvers actually caching things? Sometimes we think they are. Sometimes we think they may not or they may cache for a shorter time or they may cache for a longer time. So we think we can get some measurement of that.

We also can look at some of the DNSSEC and IPv6 distribution to figure out if the resolver is configured with DNSSEC or not or if it's capable of using IPv6 or not. We could also potentially find the most used resolvers. The most used resolvers I should qualify by eyeballs because the system is relying on Google Ads, so this is used by physical users and not machines. So it's not

going to capture machine-to-machine communication but user-to-machine communication.

This might be some measurement [inaudible] project that could also educate us about the key rollover and how many resolvers are actually really needed to cover 95% or whatever percentage of the population we would like to.

This is a work in progress. [Those] are new metrics that Geoff would like to propose. In the same spirit that Christian described earlier, we want to make this automatic so we can collect the measurement and track this over time for several years.

So in a nutshell, that's the project we would like to do with Geoff.

I apologize for him not being here, but there must have been some outside circumstances.

If there are no further questions, then we will just close this session early. Oh, question now.

RUBENS KUHL: Actually, it's more of a response to Pat Kane's comment. Why are the number of reporting DNSSEC resolvers have increased the number of those reporting [inaudible] 2010 KSK? We do not yet whether those are validating resolvers or not. So this could

be someone that indeed only has the root key but is not validating. So it's not a possible problem when the key rolls. So if we do any study like that in a metric, we should probably look into validating resolvers with old keys, not only resolvers reporting old keys. Because that's not something that measures anything that could predict what will happen when we roll the root key.

ALAIN DURAND: That's a very good point, but I will add to this. We should somehow weigh this by the number of users that are behind that resolver. If it's only something that's used in somebody's basement and is being turned on for five minutes, it may not have the same importance as a resolver that serves millions of customers.

RUBENS KUHL: Agreed.

ALAIN DURAND: So if no further comments, then we will close this session. Next ICANN meeting is a policy meeting, so there will not be any technical sessions so we will not meet. But we will see you all in Barcelona.

[END OF TRANSCRIPTION]