BARCELONA – Identifier Technology Health Indicators
Wednesday, October 24, 2018 – 10:30 to 12:00 CEST
ICANN63 | Barcelona, Spain

UNIDENTIFIED MALE: ICANN 63, Wednesday, October 24th, the Identifier Technology Health Indicators.

CATHY PETERSEN: Good morning again, everyone. Welcome to the Identifier Technology Health Indicators session to be lead by Alain Durand, Principal Technologist of the Office of the Chief Technology Officer at ICANN. Alain, please go ahead.

ALAIN DURAND: Good morning. I have unfortunately lost my voice today, so I will try to still speak, but I will let my other panelist here do most of the talking. So this is the ITHI session, so ITHI is Identifier Technology Health Indicator project. This is a project that we run from the office of the CTO, and the idea is to track over a long period of time the "health" of the identifier system.

So we spent a couple of years trying to define some problem areas that we thought were important to monitor and then define what could be measured, and now we're in the phase where we are gathering data to actually compute those metrics.

So there are two different phases in this, or two different, I will say, sister projects in ITHI. One is run by the IR for the number part of the house, and one is run directly in our Office of the CTO for the name part of the house.

So Paul Wilson, who is the chair of the NRO, will make a report on the state of the project in the IR area, and then I will let Christian Huitema, who is working on this project for us, to talk about what is happening in the name space. And then we have some discussion about the new ideas on future directions that have been coming up recently.

So without much introduction, I will leave it now to Paul to give us a report on what is happening in the number space.

PAUL WILSON: Thank you, Alain. Hi, everyone. Pleasure to be here. Thanks for coming. I'm the head of APNIC, the IP address registry for Asia Pacific. We are one of the five RIRs that are coordinated under the banner of the NRO, the Number Resource Organization.

What I'm reporting on here is a joint project of the Number Resource Organization which has involved all of the five RIRs in responding to this suggestion from ICANN to join in a project called ITHI, which is looking at sort of aggregating or at least taking a consistent approach to some concepts of health of identifiers under the ICANN banner, which include, of course, the IP addressing number resources.

So the RIRs, as you probably know, are responsible for allocating IPv4 addresses, IPv6 addresses, autonomous system numbers, and for

running the WHOIS databases that provide the public registration information for those resources.

So we are authoritative for that set of identifiers and we're quite happy to come along with ICANN on this project which is – it's very similar to and it sort of encapsulates or it relates to quite a number of priorities which are already well-established in the RIR world, namely to do with the correctness and the completeness of our WHOIS databases and to make sure that those services are able to provide the service that they're expected to by way of having the right information.

So we haven't expressed that in terms of identifier health, we've more spoken about the effectiveness, correctness, completeness of the registry. But we've sort of adopted this terminology in the context of this project.

As we entered into it, we felt that the RIRs ourselves, we're sort of an autonomous, complete set of communities serving stakeholders who are registrants in the RIR databases or who are users and kind of relying parties, I guess, on the RIR databases.

So we started by actually not trying to impose our own view on the communities in terms of what we'd be spending resources and time on here, but we started, as we often do, with a survey of our stakeholder communities to find out what were the important indicators of health of – as it's referred to, health of IP addresses and related number resources.

And the point of that – can I have the next slide, please? Is there a clicker here, or is it remote controlled? Okay, thanks. So, we started with a survey which I'll talk about here and I'll point you to the documentation. But from our point of view, what is ITHI? It's a project initiated by ICANN which is part of the goal of improving security, stability and resilience of the Internet's identifier systems, and specifically, developing metrics to measure that health as it's defined. And as I've said, we've agreed to come in on that.

We've got, under the Number Resource Organization, it's a coordinating body and it coordinates across a few dimensions of RIR activities. So we've got a Registration Services Coordination Group, this is a group of staff of each of the five RIRs who are involved with the registration services, and the Engineering Coordination Group, and obviously, similarly, that's a group of staff responsible for the infrastructure of the servers and services that we offer.

And so it was actually those two groups that came together to develop the initial set of metrics which were then put out for sort of Public comment. Next slide, please.

And that started in towards the end of 2016, and we spent much of this year on developing that set of indicators, conducting the community consultation, getting the feedback, of which I have to say there was not very much. So we were largely going on the basis of what our original proposal was. There was a final draft submitted or completed in May this year, and we've been working on actually – on working out the

common specifications that will allow us to fulfill the goal of measuring these indicators.

The RIRs are independent organizations, and we've got similar but different systems, so there's quite a bit of engineering coordination and registration services coordination to make sure that what we're doing is talking about the same types of objects in our respective databases, the same terminology, and that then on the engineering side, that we're producing the results in a consistent format.

So it's been quite a bit of work across five RIRs during the course of this year. And we're in the process now of validating the measurement specification from each of the RIRs with a goal currently still of starting publication on the 1st January next year. There's an obscured URL there, which is www.nro.net. That is the NRO's website, and that would be the publication point for these results when they're published. Next slide, please.

The scope of the metrics for the Internet number registries, it's basically, as I said, it's all of the Internet number resources which are administered by the RIRs, and those are published currently in a couple of different ways, both in WHOIS databases and also in a consistent what we call a stats file format, [known as] the extended delegated stats files, and that's a consistent view of all of the IPv4, IPv6 and ASN blocks that are under RIR administration that have been delegated by the RIRs to other parties.

So the scope of our measurement are the records in the first place that represent direct delegations from the RIRs to parties of those resources.

So we're not talking about any other blocks, whether V4, V6 or ASNs that might be reserved by the IANA, future delegations, special purposes, etc.

So part of what we went to in our proposal to the community was that we were defining accurate data as comprehensive being complete and unique, being correct and being current. So those are three sets of categories under which we're performing another set of specific measurements. Next slide, please. Next. Next. That one, yeah.

So the first of those measurements was comprehensive, being complete, and that simply means that all of the Internet number resources administered by each of the RIRs are accounted for. It's a matching process between the IANA registry in which the allocations to the RIR's recorded, and as I said, those delegated stats files which they see in the raw format on the right-hand side, there's simply a matching process between those two registries to make sure that every address block delegated by the IANA is properly covered by corresponding records in the delegated stats files. Next.

Unique – sorry, this is not uniqueness yet. This is completeness of the data, which is that registration data is complete for all of the INRs which are accounted for, and that's a matching process between that representation of the registry and the stats file on the left, and the registration records, of which you see a sample on the right, and that's a WHOIS record in the RIPE format. And we need to make sure that we've got a contact and legal name, a complete registration for each of those Internet number resources under our management. Next.

Uniqueness is a fundamental goal of the RIRs. I would hope that every report of non-uniqueness that might happen is a very rare thing, if it ever happens at all. But we do have to make sure, and part of this is naturally to make sure that every INR administered by each of the RIRs is uniquely registered and uniquely administered. And of course, to detect any problems.

So again, this is something that is naturally part of our role and responsibility which we've been validating routinely for many years now. But as part of the ITHI, it's also one of those health indicators. As I say, it's an indication of very poor health if the RIRs are not able to uniquely register IP number resources. Next slide.

The next measures are correct and current, so we want to make sure that Internet number resource registrations match official sources. There are regional policies and procedures which actually do match up the holders of Internet number resources with legal persons, legal entities in each of the RIRs, and that's also definitely a health requirement.

And not only that, but increasingly, we're making sure that those resources relate to functioning and reachable parties. So it's not enough that there be an e-mail address that doesn't bounce, but all of the RIRs are making sure that there are relevant, reachable parties at the end of those e-mail addresses who will respond when they receive correspondence. Next, please.

So what's next, that's a high-level set of the indicators that we're dealing with. At the moment, each of the registration services teams

and the respective RIRs are working with engineering teams to validate the data that's being measured. Those engineering teams are implementing, finalizing [how – the measurements we take and implement it,] and that will determine the date when the first results will be completed. But as I said, the aim is 1st of January next year. Next, please.

There's a website URL that I'll give you on the next slide, but there is a lot more detail, of course, on the website in terms of how we exactly break down these categories of health metrics which I've mentioned here, how exactly they're going to be reported, what are the metrics that are reported and so on. So here's an example of some of that information, again, breaking down comprehensive, correct and current into the sort of subtypes of metrics that we've got.

And on the next slide – please – there is a couple of URLs, one for the nro.net/ithi/project which is where you'll find more of this information, and also [a pointer] to the ICANN ITHI FAQ. And that's all I have to report from the RIRs at this point.

ALAIN DURAND: Thank you very much, Paul I think this is really a lot of work that has been put into this, and we're quite happy [now] to see the results and data coming in. We are putting the data we collect in the name space into the ODI project as it was presented, I think, yesterday. Do you think it will be possible to find a way to export your data when it will be available into the same pipeline through ODI?

PAUL WILSON: I'm sure it's possible. I'm not sure that we've seen or if any of the teams have actually interacted with the ODI in terms of looking at what your metadata schemas are and so forth that we'd presumably need to adapt to to provide that data. But of course, technically possible. Thanks.

ALAIN DURAND: So, we would love to start a conversation on how to do this.

PAUL WILSON: Okay. We'll refer that to our engineering group to see where they are exactly on that, on looking at ODI. Thanks.

ALAIN DURAND: Thank you again, and it's really great to see progress being made there. I'm really happy. Is there any question for Paul? Maybe in the chat room?

CATHY PETERSEN: No.

ALAIN DURAND: Okay, so if there is no question, then we will move on to our next segment. Christian Huitema is going to talk about the name side of the project.

CHRISTIAN HUITEMA: Good morning. I'm Christian Huitema. For the last year I've been developing with Alain this set of – this measurement system that we have to check the health of the name-oriented identifier technologies, and – can I get the next slide, please?

So ITHI, we started the project, and from the beginning of the project, we had a couple of big principles that are already started when I was doing this presentation in Puerto Rico six months ago, but basically, we are a technical project, we are not trying to interpret the data. We are strictly about measuring what's happening, and we are leaving all the interpretation to somebody else.

What we want to have is a set of measurements that are tracking problem areas, and we want to operate on a long duration. All the methods that we are doing are meant to be automated as much as possible, providing measurements day after day, month after month, so that we can have this long-term service and five years from now, we can compare what's happening in 2018 and what's happening in 2022, 2023 and see all the trends and see that.

We had a presentation of ODI yesterday. All the data that we are producing are sent over ODI for people to interpret. So, what do we have? Can I have the next slide, please? We have eight areas of measurements that we are doing now. We have the first two areas about, I would say, the quality of the data and the quality of the registration data.

First, M1, [to ask] the quality of the WHOIS data, and we are getting that. We are not doing the measurement ourselves, we are getting the data from ICANN's compliance department. Our added value there is to take that data, get some significant statistics out of it, and have those statistics copied in our database and made available over time.

We do the same thing for the domain name abuse data that come from the DAAR project and give us idea of how we see evolving, say, the number of spam domains compared to the total number of domains, and we look at that over time. And again, the idea is to keep data so we can come back and look at evolution and look at it in a public way.

The next metrics are about the DNS system itself. The M3 metric is about the root of the DNS, and we are looking at traffic at the root in order to find trends and detect, for example, the rate of success of queries, the amount of [bizarre names that we are looking,] name leakage, these kinds of things.

M4, that's the same thing at the level of the recursive, so we are instrumenting a set of recursive servers all over the world and we are looking at the traffic that they receive, and we are characterizing the traffic to understand for example what kind of names people are looking at, what kind of usage they are doing with DNSSEC. We are doing all that at the recursive.

Then we have analysis of the resolver behavior. We want to check how resolvers are processing to resolve names, whether they use DNSSEC for example or how they manage caching. We are looking at the IANA registries that are linked to the DNS to see whether they are healthy,

and we define health as, are people using the values that have been registered? Are people squatting on unregistered values? And we are looking at that for all the registries that are related to the DNS.

We have a specific metric on DNSSEC deployment at the top-level domains, and then we have analysis of traffic at top-level domain servers to basically, again, do the kind of thing we do at the root. Look at the traffic, look at usage patterns and detect that. Next slide, please.

A key point of what we are doing is to ensure that we can have distributed measurement systems. We really on partners all over the world, and the big issue there is to maintain basically the security and the privacy of the data. The way we do that is by not running any kind of measurement network ourselves. When we want to extract measurement from, say, a recursive server, we are working with the operator of that recursive server. We are giving them a piece of software they can run, but they run it on their machines. That's what we have to process the raw data. We never see the raw data.

What the software does is take the raw data, extract statistics which are anonymous, have no PII whatsoever, and they are also much smaller than the original data, and then we have this kind of wall between the operators and our servers. The operators will have a login on our staging server so that they can upload data. We do not have any kind of logging on their servers, it's strictly separated.

Now, when we upload the data on the ITHI staging server, we do analysis of the data there, we prepare the extraction of the metrics, we put that in two metric files, and the metric files are then published by

the Open Data Initiative so that anybody can then look at them. So we have a process that is respectful of privacy and provides data for researchers. Next slide, please.

Currently, we have a set of initial partners. We have just started, but we are grateful to have these initial partners. The university of La Plata in Argentina, the University of Cape Coast in Ghana, the Nawala DNS operation in Indonesia, and the Kaznic server in Kazakhstan. We have several more partners that have told us that they want to work with us and they will be coming online in the next months. Our goal is to expand this system and we are making a lot of investment to make it easy for people to partner with us. Please.

I did a presentation already of the ITHI metrics back in Puerto Rico in February. Since then, we have developed our system and we have added new measurements. The first thing we did was to partner with APNIC – and Geoff Huston is on stage with me here – and to measure the integrity of the resolver service.

APNIC has an interesting system in which they can send probes to a large number of customers all over the Internet, all over the world, and then look at the way those probes – they send DNS queries from those probes that [are targeting] their own servers, and by doing that, they can analyze the way those probes are processed by the DNS resolves, and that gives us an insight into the way the resolvers manage to resolve queries for the DNS.

And we look at things like cache management, do resolvers refresh their cache as expected in the DNS, in the TTL? We look at things like auto

refresh. Do resolvers proactively refresh their cache? We look at the number of resolvers that attempt to do DNSSEC or at least signal that they might attempt to do DNSSEC by putting up the DNSSEC okay bit in their queries. And we also look at resolvers that actually perform DNSSEC.

If you control your server, it's relatively easy to sometimes serve a file that on purpose has invalid DNSSEC data. And if the DNSSEC data is invalid, the client should not be able to resolve the name. So we see that, we see how many resolvers will actually reject those invalid data. And to disclose data, it's about 25% right now.

And we look also at concentration – so what we are looking at today is the fraction of clients that are served by the top 10,000 resolvers. And we'll look at that. We can look at the original data that we have that's the –

CATHY PETERSEN: Next one?

CHRISTIAN HUITEMA: Yes, please. The first batch of data – all those data are published on our website, so you can look at them in real time if you want. Okay, but basically, for that metric, we have a set of sub-metrics. Each of them gives us a fraction, like for example here, I see that 27.9% of resolvers will refresh the data before the cache expires. Or I can see that 84% of resolvers set the DO bit in queries. And I can also see that setting the DO bits is not quite the same thing as actually doing DNSSEC, because only

about 25% of resolvers in fact do DNSSEC and reject a name resolution if the DNSSEC data is invalid. So that's a new system that came online in July and is online now. Next slide, please.

The other thing that we did since February in terms of progress of the project is to look at software updates. We have realized that it was a major problem for the operators of the resolvers to take our software, because our original idea was to make the software available on GitHub.

And it is still available on GitHub, everybody can have the source, can inspect the software, verify that we are not doing something unhealthy with the data or with the code, and the package is available for Windows, Linux, FreeBSD. I mean we verify it compiles. But we los have packages that are precompiled that are available for – [to] release of Linux, Centos and Ubuntu so that people can just do their usual sudo apt-get install or whatever it is to get the last version of the software and have it in real time.

And that allows us to evolve the software over time, [inaudible] fix bugs – not that we have any bugs, but we might want to fix them anyhow – or to add new measurements if we have a demand from the community to add another measurement in the system.

The software has been audited. We had our colleagues at NLLabs run a complete audit of the software to verify that it was done proficiently and that there was nothing bad. And they were fairly positive. They added a couple of things that – asking us to do automatic checks and things like that, which we did. But effectively, that gives you some

guarantee that when you're using that software, that software has been audited, it's good quality, it's not going to create a security issue in your site. Next slide, please.

Finally, the last metric that came online was a metric on authoritative servers. [inaudible] the first authoritative server that we have is Kaznic in Kazakhstan, and we are measuring a number of things on those servers.

We're measuring whether they are failing queries and what's the fraction of queries that are failing, which surprisingly is quite high. We are measuring the behavior of the resolver, what are the use EDNS and which EDNS options they do use. We are measuring whether they said that they are using DNSSEC and we see that number of 82%, which is very similar to the number we are seeing with the APNIC system, so that gives us a way to corroborate the data and verify that.

And we have also started to measure whether the resolvers are doing QName minimization, which is privacy techniques that have been developed by the IETF to restrict the available data at the resolvers and avoid metadata leakage.

So that has been online since this month, it just came out, and we are very happy with that and we are recruiting more partners to get more data sources and have higher quality data. Okay.

Couple more updates. We have been updating our metrics, our measurement at the root and at the recursive to track DNSSEC and track QName minimization. We are now publishing the data on the

ICANN website at ithi.research.icann.org, and it's completely available, anybody can look at it. and we have started integration with the Open Data Initiative publishing the data as we saw yesterday during the ODI session. So in a word, we have been very busy. Next.

Now, at that point, when we have all those data available, we have to take a step back and say, "Okay, what do we want to do exactly?" And that was a question, a challenge that we got that says, "Okay, we have a rich set of measurements, and they're a value by themselves," but the original mission was not to have a whole lot of measurements about the DNS. The original mission was to say, "Hey, are we in a healthy ecosystem or not?"

And so our next stage now is to go from the set of measurements that we have already started to acquire to some kind of a health dashboard for the identifier ecosystem. We have rich metrics, but they are a bit too rich, so we want to have something like the executive summary of the system. If you want to have a single page to tell you, "Are we doing fine?" What will be on that page? What will be the dashboard of the identifier technologies?

We are going to put that on the main page of our site. It's not done yet, but I want to give you a preview today. Yes, [getting there.] So we are going to have four indicators. We have looked at the system, and I think there are four things that [pop out.] And we'll see them one by one in the next slides. Please.

Yes. The first is looking at the traffic at the root. And this is an extract of our M3 metric. And what we see is that traffic at the root is largely based

on non-existing domain names. Like about two thirds of the traffic are requests for what is the address of this or that name which does not in fact exist as a registered TLD.

Only one third of traffic is actual requests for actual names that exist. And in fact, even within that third of the traffic that is for reasonable requests, we observed that the larger fraction of it are repeated requests of resolvers that repeat the same request again and again even so the result less than a few minutes ago. And the TTL is [inaudible] typically.

So we have this system there, and what we see is that – I mean if I look at a graph like that, [inaudible] not really good, we see that. And it's probably sustainable. It's not evolving in a dramatic fashion towards getting worse and worse, but it's still not very good. So we think that that's something we want to show, and for me, the indication of health is if the dark gray area there shrank over time, which it is not doing.

And there are ways to do that by having better actions at the resolver, better filtering. Also working with various actors to stop the business of using fake names in the DNS and things like that. So that's the first indicator that we want to pop up at a high level in a summary of the system. Next slide, please.

Geoff is going to speak about this concentration system next. It's not something that we are actually measuring today, but it's something that we observe as kind of metadata when we are doing the analysis of resolver behavior.

In the study, we had about 28,000 resolvers providing responses when we do these millions of probes over the Internet. That means that that's a good indication of the population of recursive resolvers which are actually deployed today. What we see is that this distribution is very skewed. There are a couple of big resolvers that have a massive number of customers, and a very long [tail of very long resolvers are] very few of them.

We believe that this concentration can potentially be a measure issue for the health of the ecosystem, because if you have concentration of this resolution to a very small number of actors, those actors then create some kind of a point of fragility in the system. So we want to track that, and Geoff will discuss that after my talk. Next.

We also want to track DNSSEC deployment. It's kind of a priority of the current DNS management to increase the security, and our security tool is clearly DNSSEC. We have a number of metrics that are available today by our systems. I listed them all on the slide. Measuring the fraction of TLDs that are signed, measuring the number of resolvers that attempt to use DNSSEC, number of clients.

In fact, I underlined here two important metrics for me. The number of resolvers that actually use DNSSEC, which is about 25% today, and the number of stub clients that say that they might use DNSSEC by setting the DO bit in the query, which is a bit less than 1% today.

Those two numbers should give us pause. First, 25% is way better than zero. That means that DNSSEC is actually used, and since those

resolvers doing the verification tend to be big resolvers, the fraction of clients that might benefit of that is fairly large.

So that's a good indication. We'd like to see that number, 25%, grow towards 90-99%. That will be basically the good result. Then next thing we'd like to see is that having the clients also verify the data. And today, a very small fraction does.

So those two numbers, 25% and 1% - 25% that will be, if I say, in the yellow area, 1% will be in kind of the reddish area, means that basically, customers kind of benefit form DNSSEC because they might use a recursive resolver that may be checking DNSSEC, but they are still exposed to all kind of shenanigans in the traffic between that recursive resolver and themselves. Having that grow will be an indication of success in the deployment of DNSSEC. So we want to pop that, again, in our dashboard as one of the key indicators. Next slide, please.

And finally, the fourth indicator that we want is the name leakage. I have summarized here the result of two metrics, M3 and M4, measuring the leakage of names at the root and at recursive resolvers. Basically, there is no surprise there. the main two names that pop up are dot-local and dot-home, which are leaked by a number of home networking systems.

If you look at the list, there is no big surprise. These are the names that you would expect. You will see dot-corp. You might see dot-mail somewhere down the list, but it's a bit smaller so it's not in the top 20 that I've put here.

And you'd think that we'd see the same kind of names at recursives. There is some correlation. Dot-local and dot-home are indeed at the top of the leaks at resolvers as well, but there are also not so much correlation because the leaks at the local resolvers tend to be dominated by local issues.

If you look at the data, you see effectively a mix of two things. Clients that attempt to use names that are not actually registered, and misconfigured devices like Wi-If routers that try to discover each other using names that are configured by the manufacturers.

Because we have a small number of recursive, those local elements today dominate what we see at the recursives. That would basically get better if we have more and more data sources, because at that point, we'll get an averaging.

Anyhow, this gives us a good indication of what are the names that are currently leaking at the root, what are the significant names. And I think that's a good data point for people making decisions about the DNS. So these are our four indicators.

The project is ongoing. We are recruiting more partners. We are publishing our data, we are constantly verifying that we don't have any kind of analysis. We are also getting feedback from the community about new indications that we might want to track. For example, one piece of feedback I got this week is that it might be useful to also track internationalized domain names and how they are used, and we could do that by revising our software and adding probes for that, and it will be interesting.

But as a general practice, we want to have feedback from the community about what else would people want to see, what is interesting, what is not. And we really need partners, so please contact us, contact Alain or contact me and we will help you set up a probe in your system and you will get all kinds of access to data that might be interesting for you as well. Thank you very much.

ALAIN DURAND: Thank you, Christian. Sometimes, service provider of top-level domain operator ask me, "What is in for us? Okay, it's nice we collaborate to worldwide project, we share data, but what is in for us?" And I have two answers for that.

First one is it enables you to compare your network from a DNS perspective to other people's network and see if you're in the same ballpark or if you have some differences. And if you have some differences, you can try to understand why, and maybe some good indication.

The second thing is it can be used as an early warning system. Very often when there is an attack, in the weeks leading to an attack, you may start to see different DNS pattern, different DNS queries when attacker is actually probing your network, probing your defenses.

If you have a baseline that you run every day of what your DNS traffic looks like and you start to see that you're deviating from this baseline, something might be happening. It may be something you want to pass

on to your security team. So this is something that could be another motivation for people to collaborate to this project.

But before I take questions for Christian, I would like to pass it on first to Geoff, and then we will take questions for all the panel.

GEOFF HUSTON: Thanks, Alain. My name is Geoff Huston, I work with APNIC. Can we have that slide back again, please? Wherever it went to. Right, thank you. Up the top first. You're too low. Go back a bit. Okay, so let's speak to the first half of the third slide. With luck, we might even see all of it.

The last ten years on the Internet have certainly had some disturbing revelations where the area of surveillance, which used to be a relatively dimly lit world, was exposed to a very harsh light of leak documents and their repercussions. It became evident that the DNS is widely used and abused as part of that process and that your queries in the DNS are an accurate and up to date reflection on what you're doing.

It is also the case that in terms of access to content and regulating or controlling that access to content, the DNS is seen as one of the most commonly used points of control, that by preventing certain names from being resolved by the DNS, users in theory then cannot get to those corresponding services.

The IETF in a set of meetings around five years ago took the stance that this form of – one would call it abuse, perhaps, of the DNS – was a form of attack on the protocol. And they then engaged on a concerted program in the DNS to introduce a number of elements that surround

your queries with a cloak of encryption. And we're now seeing three of these gain some degree of public prominence as that work now comes out into deployed infrastructure.

On the web these days, we are very used to seeing that standard lock icon and things going green in in the browser bar. In other words, an indication that when you are entering data into a web screen, you are talking directly through an encrypted session with the server you intended to talk with. And in theory at any rate, the encryption is such that there's no third party that can eavesdrop or see what you're saying.

This is called transport-level security, or TLS, and is now very commonplace in the web, but in the DNS was very, very rare. Your queries were in the open, anyone who cared to look could see, even over Wi-Fi.

So the first of these is a toolset called DNS over TLS. It uses an encrypted TCP session directly from the end client, your system, to some chosen recursive resolver. It's available now on the Android Pie release, and you'll see it also coming out in a number of browsers over the coming months as they introduce this.

Again, the whole idea is anything you say in the DNS will be encrypted on that first hop across to the recursive resolver. But HTTPS uses TLS, and one thought has been, why don't we just put the DNS into the web itself and make the entire DNS transaction look a little bit like an HTTP get? Instead of getting a webpage, you would get a DNS response. Why is someone moving those slides? Leave it there. Thank you.

So DNS over HTTPS uses that same level of TLS encryption, but it can now be used not only by your system but by an individual application within your system to lift itself out and undertake queries directly against a recursive resolver with no other party being able to see.

Google originally started [on] a separate protocol called QUIC which is now being taken up by the IETF. In the first two forms, TLS and HTTPS, the TCP control part of the session is still in the open. And there is a lot of middleware out there that manipulates, alters and otherwise mangles that TCP session.

QUIC is the response to this, where the visible part of the session is simply a UDP packet stream, and everything else is hidden inside the encryption layer. The transport protocol you're using is now only known to you and the party you're talking to, the network can't intervene.

Now, the side effect of using these in the DNS is quite curious, because now instead of using your local DNS resolver that you got from your ISP or whatever you configured [in, if you will,] automatically, so if you bring up a mobile phone, DNS is pre-canned, once you start using these applications and these browsers and these tools, you can redirect the queries that you're sending out across the local infrastructure by passing it to anyone you want.

Now, is this good or bad? And one of the questions this raises is, to what extent do a small number of players tend to be dominant in the DNS recursive resolver world? So now let's flip to the next slide. Thank you.

Resolver centralization. Oh, that's too big. Let's leave it at that. Thank you very much.

So let's take a simple statistic. As we see from the APNIC experiments and looking at DNSSEC, we also started looking at to what extent users out there use Google's 8.8.8.8 servers. Quite surprisingly, we found that number to be around 14 to 15%. In other words, across the entire planet, one in seven users sends their queries to Google. They don't send it to their local ISP – or they might as well, but Google gets to see an enormous amount of DNS traffic.

This prompts the question, because Google is not the only player in the space, there are a number of others, and you may have heard of some of these. Cloudflare have recently started a project with 1.1.1.1, PCH and a few partners have started a project on 9.9.9.9. Get the pattern? And a break from this, VeriSign has been running 64.6.64.6. Highly memorable, isn't it? And there are a number of others, including Yandex, OpenDNS, DNS.WATCH, Comodo and so on. That's a lot.

So the question is, to what extent will those players start to dominate the DNS world? Because if they do, the kind of frightening prospect is if they refuse to serve an answer for a name, that name doesn't exist, irrespective of what the DNS might otherwise say.

Will this sort of adoption of DNS and the centralization lead to some deeper issues with the DNS? We don't know. But we would like to understand to what extent the DNS is centralizing and amassing. Is that share by Google getting smaller or larger? Will Cloudflare rise

inexorably and capture more users? Will we see OpenDNS, Quad9, etc. also assume a much more prominent role?

So without making a value judgment as to whether that's good or bad, it would be very good to understand to what extent this is happening. So we certainly would like to measure inside this project a project called resolver centralization to understand from a massive set of measurements every day to what extent users are moving towards using these kinds of services that hop over the immediate infrastructure and direct their queries directly to some of these larger open DNS resolver factories.

That's all I had, Alain. Thank you.

ALAIN DURAND: Thank you, Geoff. So, there was a slide that Christian presented earlier that had some numbers, the measurements from one experiment we ran in July. And just want to read those numbers again to give some context to this. The largest resolver we saw [had there] about 13% of market share, meaning 13% of eyeballs that were touched by this experiment were using this top-level resolver.

It took only 25 resolvers to account for 50% of the eyeballs. 25. If we want to go to 90% of the eyeballs, we only need 460 resolvers. This out of about 28,000 resolvers we saw that day. So 460 resolvers were enough to account for 90% of the eyeballs. This is just one measurement, and we're suggesting essentially to refine the system because defining what is a resolver is sometimes a tricky question, and

to keep measuring this over time to have an indication of, is this concentration that we are seeing getting worse, or are we seeing some diversifications?

Now, [inaudible] like to open the floor for questions, and please identify yourself for the microphone and the record when asking questions.

ALISSA COOPER: So I have a couple of questions, Christian, on your presentation. But actually, maybe this gets to the point that Alain was just making. So the statistic that you showed, Christian, I think under the M5 category, was related to the top 10,000 resolvers, but then as Alain just reported, there's some further detail that you provided about different cuts and you looked at the top one and the 25 resolvers serving 50% of clients. So I was just curious in terms of the data that's available to the rest of the world to go crunch and look at this. Are the buckets defined by you, or would anybody be able to look at concentration at a variety of levels based on the data that you've made available?

ALAIN DURAND: So let me answer this one. The buckets that are used for this particular experimentation is something that we define, because it was [one of] experimentation. Now, if we [run] this as an ongoing measurement, then we will publish probably the full set of data. But I will have to check with Geoff if it's something that will be okay to do, of course.

CHRISTIAN HUITEMA: Let Geoff answer, because he's the one who's collecting the data.

GEOFF HUSTON: We've certainly had our own internal issues about whether resolver IP addresses sit naturally within personally identifying information or not. Some consideration of this has kind of lead us to as long as we are never talking about who asked, the resolver addresses themselves, particularly on the larger ones, do not appear to reveal any particular information about you, me or anyone else.

And so we're certainly amenable to releasing the data set of the larger ones. But once we start to get to ones and twos, I think all of us get a little nervous. So in looking at this larger set of resolvers and centralization, we are certainly happy to release some of that.

There is another caveat though about when we talk about resolvers against users, because two things happen. When you get a large public system like Cloudflare or Google or Quad9s, they don't actually use those addresses on the resolution side. So there's a second mapping set of the addresses used by each of those resolvers.

That is discoverable, but not necessarily published. And again, I think some further thought on our side as to whether the resolver operators are comfortable. If we discover the mapping information and release it has yet to be determined. So even if you see an address, it's not clear that that's an address used by a public resolver or not.

The second part of this is counting users against resolvers. We have no intention of releasing user counts. It gets a bit close. We're happy to talk

about percentages where the absolute scale is no longer there. So with the caveat around grouping IP addresses to resolver services and understanding the distinction between percentages and absolute counts of users, we certainly intend to release some of this data, and it would allow the measurements to be replicated by others, we believe. So hopefully, that would be adequate. Thank you.

ALISSA COOPER: I'm just thinking out loud, but isn't the volume of queries potentially almost as useful as users? No? I mean in terms of understanding this concentration metric and doesn't have the same privacy effects.

GEOFF HUSTON: We have done some work – and then I'll hand it over to Christian – there are kind of three kinds of resolvers out there. There's the sane ones that seem to ask real questions. There's the replayers and shadowers that seem to ask old questions at moderate rates, and then there are the braindead loony headbangers who seem to occupy around up to 60% of resolver queries asking the same question in a demented fashion. So when you say, "Are query volumes useful?" The headbangers dominate everything and make the numbers totally weird. Christian?

CHRISTIAN HUITEMA: Yeah. One thing that I've been trying to do is make sure that we have several sources of data, several ways of looking at the data for each of those phenomena. For example, when we were looking at the DNSSEC deployment data, I could corroborate the data that Geoff is providing

with his APNIC experiment with data that we are measuring by looking at the traffic at authoritative resolvers.

And by having several views on the same phenomenon, I think that we can eliminate a large bit of the bias and make sure that we are not looking at faces in the clouds or whatever. So I think that's a key characteristic of the project. [I mean] look at that in several ways using completely different methodologies so that we can corroborate the data.

ALISSA COOPER: Thank you. That's very helpful. Just two more quick things. I do think for some of the statistics that you showed, having some understanding of the volume of queries over time could be useful. So for example the NX domain results that you showed, 68% of what hits the root, if the volume of queries over time went down dramatically, that would be an interesting other data point along with the percentage of junk queries, essentially.

Overall, I think at least from the IETF and IAB perspective, understanding how the resolver concentration is changing over time would be very valuable, I think. It's something that we've been talking about, so I just wanted to express my support for that.

I do think it would be useful to distinguish between or maybe collect separate statistics around which transport protocol is being used versus the concentration. Like it's possible that everybody goes to 8.8.8.8 even if nobody ends up using DO or [DOT] or DNS over QUIC. So

just to maintain a little bit of that separation conceptually that we don't necessarily conclude that the evolution of the transport protocols means something for concentration when they could be having separate effects. I don't really expect that, but I think it's important to distinguish those effects.

ALAIN DURAND: So a question for Geoff. In your system, do you think we can actually measure if people are using this DOH or DOT or QUIC?

GEOFF HUSTON: Thanks for the question. Look, the DNS is a hideously ugly beast. It resists most forms of measurement. It is completely opaque. Questions go in, queries to authoritatives come out. If you're lucky, they're related. 40% of the time, they're not. No matter what protocol the end user might use to connect into their first hop recursive resolver has almost no bearing on whether the consequent set of [inaudible] and the query to the authoritative pops out at some other point in the cloud.

So certain measurements make certain aspects of the DNS clear, and others don't. In the measurement that we do in APNIC, we seed unique questions to end users via scripts. So we understand they've generated a question and they've sent it to their recursive resolver. Can't see that happening though because we can't see that.

What we do see is the consequent [inaudible] query coming to our authoritative. But the protocol they're using is the protocol between the recursive and the authoritative, nothing to do with the user. So it's

much harder to see what users are doing. And for that, you need to move into recursive resolvers. And I'm going to leave that to others because we certainly do not see or measure what recursive resolvers do.

The only thing I'll add just before I stop though is the one thing we do do in our DNS names is add the time of day when the name is created. So what we understand when we look at our authoritative servers is the age of the name as it hits us. And disturbingly, between 40 to 45% of all the queries that we see at authoritatives are for names that should never be asked about because they're old and dead.

And that sort of is an observation of the broader DNS, that all it seems to be doing is repeating up to around 40% of junk. If you really want to impact upon the query rate in the DNS, eliminating that junk might be a wonderful thing, we just don't know how it's being generated. Thank you.

CHRISTIAN HUITEMA: Well, actually, we know for some of it. I can point you to the ten lines of code in the Google Chrome code that generates funky names in order to test whether they are connected to [inaudible] or not. But there's some of that.

but Alissa, to your question, there is a problem there, there is a huge tension between measurement and encryption, because a lot of the measurements that we are doing are by using effectively Ethercap. Excuse me, Wireshark.

Basically, we're looking at traffic and we're analyzing the traffic as it passes. Now, it's very clear that it's a challenge for us that if a larger portion of traffic was encrypted, we would not see the data. That doesn't mean it's unsolvable, but it means that it's on a workplan that we have to evolve our probes over time so that instead of attaching directly to the ethernet and looking at the packets, they attach inside the software of the resolvers and look at what the resolver actually sees after decrypting the data.

And there are ways of doing that, and clearly, measuring the fraction of traffic that is encrypted and measuring which fraction comes over DNS, over TLS, which fraction comes over QUIC, which faction comes over HTTPS, that's clearly something we need to do, and that's something we'll try to put on the workplan and doing the instrumentation. It's a more complex instrumentation than what we have now, but we need to do it, clearly.

ALISSA COOPER: Perhaps I could clarity, which is that even if you don't do that, I think it's important then not to assume correlation. That's all.

CHRISTIAN HUITEMA: Oh, yeah.

ALISSA COOPER: So if there's things that we're not measuring, fine, but don't blame the transport protocol if it can't be [removed from the] data. That's all.

CHRISTIAN HUITEMA: No, it's clear. These are two parallel phenomena and we have to measure both.

ALISSA COOPER: Yeah. Thank you.

GABE FRIED: Hi. I want to go back to Paul's discussion about the registry databases for numbers. I know that as the IP address pool runs out and the transfer market takes off, a lot of those records are becoming updated with either new RIR members or existing RIR members getting additional addresses. But as it relates to moving IP addresses out of legacy status or out of status of not being utilized [inaudible] on the Internet, is there a health metric for address utilization in both protocols once those addresses have been delegated to an RIR member? Is that on the agenda?

PAUL WILSON: Sorry, I'm not quite sure what you're asking. In the case of transfers, transfers between registered members or customers of RIRs will be fully registered in WHOIS, both the source being updated and the recipient being updated.

There's a separate – I think every RIR as part of its transfer policy, I believe, maintains the transfer logs, so there is another dataset that allows you to check which transfers have taken place in the past. But if

you're talking about sort of what we used to call black market transfers where someone has made a deal with another party to lend or sell address space without changing the registration records, then we're not looking at that as such. That requires analysis of the routing table and sort of heuristic analysis of how addresses may have once been routed and where they may today be routed.

GABE FRIED: No, that's not my question. My question actually has to do, I think, with some of the work that Geoff has presented at other meetings, looking at the utilization of the address space in v4. Because I know that as a general health indicator, I don't know if you guys have created a metric or if anyone cares about to what extent their utilization of the v4 space is increasing over time as the world tries to adopt v6.

GEOFF HUSTON: So the metric that I did is really a huge amount of walking out on a plank. I have to take the only data I can find about the number of users in each country, and the only data that's available is data published by the ITUT's stats division.

And that's a few years old, so you take the UN stats division current population data and look at its growth rate and extrapolate across to user population data. Okay? So I now know that there are certain number of users in each country.

We then do an extrapolation based on Google's ad presentation, and this is the second leap of faith, that we kind of think that within a

country, we've structured ads so that it should be uniformly presented to all eyeballs in that country. Now, not everyone sees it, but the ad sampling rate should be roughly consistent.

The implication is that if 50% of the ads end up in ISP number one, then we're roughly saying that 50% of the users in that country are in ISP number one. And because we know the number of users – thank you, ITUT – we can do a rough correlation of users for each network. So far so good.

We then take the BGP routing data, and we look at the amount of, in v4's case, slash 32 individual addresses announced into the routing system because that ISP. And this gives us a rough – and I really do say rough – metric of the number of users per advertised address.

It has some really cloudy uncertainties, plus or minus 5%, maybe plus 10%. But it certainly points to areas of acute address stress. There are some countries where the users per IP address is incredibly high, and other evidence seems to correlate, other countries where it is not so high.

The problem with that data is really just the dramatic number of assumptions that go on behind it, and we've always been careful to try and say, "Look, this is really just a stab. It's a guess. It's pinning together a number of available pieces of data that have their own uncertainty, their own timelines," and putting them together and guessing.

Is it useful? Probably. Should we do it all the time? I'm certainly interested in feedback on that, because as I said, you really do have to

accept some caveats, and there are certain operators who would come to us and go, “We have more, we have less.” There's a lot of difference between a known reality and what we’re guessing.

But if you want us to do this on a continuous basis, we’d certainly be interesting in the feedback. Thank you.

GABE FRIED: I guess my question actually was a little bit simpler than that, which just has to do with the overall utilization rates of the v4 space in its entirety, because you have presented – and I can't remember exactly where, but sometime in the last two years – a discussion about how the transfer market is getting pools of addresses moving from unutilized to utilized, and there is some predictive value in that about when even after the last of the registries hands out their v4 space, when the last of the allocated and unused v4 space will get consumed. So that was sort of – I'm thinking about this at a much higher level, not at the national adoption rate level.

PAUL WILSON: As far as RIRs are concerned, there is a level of utilization which is very obvious and which we are interested in, and that’s whether an address block is appeased in the routing table. So someone’s received an allocation from the registry, they advertised that into the routing system via their [peers or upstream,] and by one measure, that address block is utilized, but we've got no idea without monitoring traffic, which

of course, we don't have any ability to do or interested to do or authority to do.

The next level of utilization is which individual addresses within that block are used, because of course, the block may have 1000 addresses in it, and there may only be five devices using any one of those addresses assigned – with addresses assigned to them.

On the other hand, the addresses could be fully utilized and each address could be fully used by thousands of users through a [inaudible] [NAT,] and we don't have any view of that whatsoever. I know there were research projects undertaken over history to try and work out how many individual IP addresses in the entire 4 billion IPv4 address pool are actually in play or in use, but Geoff might know a bit about that history, I guess.

GEOFF HUSTON: Yes, there were a number of research projects that tried to sort of uncover users behind [NATs] using various techniques. They generally haven't been successful insofar as the data is kind of dubious. We really have no clean way of measuring it because of the nature of the [NAT] in v4. Don't forget too that eyeball-based measurements are not device-based measurements, understanding the deployment of devices is largely a black art.

So our tools that [we know are] reliable are the BGP routing table and our allocation database. And we can map those together, but that doesn't give you what I think you're after, which is a finer-grain metric

of use of addresses, and that is something that I think we just don't understand how to measure to that level of detail.

PAUL WILSON: I'll just make one more point here which might be a bit off the topic of identifier technology, identifier technology health, and that's about the utilization of remaining IPv4 address space. Because the RIRs have got no mandate or ability to really look at the efficiency of usage of blocks which have been allocated, we assume that in the past when address space was seen to be abundant, that operators didn't care too much about how efficiently they used an address block. There could be plenty of unused addresses within a block.

We don't have any view of that. We only see whether the block is announced not the routing table and not what's actually being done with each of the individual addresses. So when the question came up about how we get a higher level of utilization of the existing IPv4 address pool, that was really one of the answers to that, was in fact the transfer, the ability to transfer addresses on the open market, because it allows an operator to make their own evaluation of how much utilization they have of their addresses, how much trouble they might need to go to to actually free up addresses and put them out on the market. But there is at least an incentive there for people to do that, and that's how the market is expected to work in increasing the overall efficiency of IPv4 utilization.

GABE FRIED: So the good news is the market actually does work in that way, because we see customers starting to more efficiently utilize [their] existing space so that they can sell [it] surplus. At current prices and as prices go up, it'll obviously become more efficient. But just as this initiative was taking root at ICANN, was there interest in the v4 transfer market or elements of it as a health indicator, or is that sort of, "Let's let that dog sleep" kind of approach?

PAUL WILSON: As a health indicator, the transfer activities didn't come up as a possible thing to be measured, as a possible metric. This is the first I've heard that it suggested it's – yeah, so whether it is or isn't interesting, I don't know. There was indeed – particularly in the days prior to the IANA exhaustion, there was as lot of interest in the disposition of the IPv4 address pool in the ICANN community, and RIR folks did turn up and talk to the GAC, talked to the ICANN board, made reports in our sessions here for the interest of the community about how the v4 exhaustion process was going, what the policies were going to be that would help to manage that, and also the transition to IPv6.

GABE FRIED: Thank you.

PAUL WILSON: Okay, thanks.

ALAIN DURAND: So, Paul, I have a question. If we get more questions like this or suggestion to get new metrics in the number space, how can we pass them to you?

PAUL WILSON: That's a good question. Normally, these discussions would come into RIRs through any of the individual RIR community processes, mailing lists and so forth. We don't have an NRO, a combined, joint internet discussion list. But I suppose actually thinking – I think [yeah, I'd lay out] here that back on that NRO ITHI project page, that is where we were consulting and asking for feedback, and there may still be an active e-mail address there for anyone to post feedback.

That consultation process is concluded, and as I said, we didn't get too much input into it. But I think everyone would see this as an ongoing evolution and certainly open to feature suggestions and requests into the future. Whether or not those can be implemented in a sort of given frame of time, I'm not sure, but I'll have a look at that maybe and open a sort of feedback channel still. Thanks.

ALAIN DURAND: Thank you, Paul. Is there another question? Or maybe we can turn to the chat room. No? Another question in the room? Going once, going twice. Well, it's lunch time so we don't want to be in your way for food, so thank you very much all for coming today. I hope it was an interesting session. The slides will be uploaded. Maybe they already are there.

CATHY PETERSEN: Yes. Slides have already been uploaded to the public schedule. The recording for this session today as well as the transcripts will be published to the public schedule also within the week or so. Thank you.

ALAIN DURAND: So thank you all very much, and bon appétit.

[END OF TRANSCRIPTION]