
DUBLIN – Root Stability Study Workshop
Tuesday, October 20, 2015 – 13:00 to 14:15 IST
ICANN54 | Dublin, Ireland

KAREN LENTZ:

Good afternoon, everyone. I think we're ready to get started. I will introduce myself. I'm Karen Lentz, Director of Operations and Policy Research at ICANN. This study is one of a number of reviews that relate to the impact of the new gTLD program, and this one obviously is focused on the impact of the program on the root server system. I'm going to give a brief introduction as to the context, and then I will turn it over to our team here.

The study stems from a preexisting commitment, among other things, to make sure that the experience of adding many more TLDs to the root and that impact on the system thoroughly reviewed and evaluated to determine whether any changes or additional steps should be undertaken prior to continuation of the program, resulting in new applications and more delegations of TLDs.

We've published earlier this year a request for proposals for qualified firms to conduct the study, received several very good proposals, and I evaluated them through the process to be able to conduct this study. We've recently engaged a consortium of

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

providers to do this, consisting of the Netherlands Organisation, SIDN, and NLnet Labs.

I will turn it over to Bart Gijsen, who will take you through the rest of the presentation.

BART GIJSEN: Thank you, Karen. Could you hand me the [inaudible] remote?

UNIDENTIFIED MALE: Yeah.

BART GIJSEN: Thank you, and welcome, everybody. As you may see on the previous sheet if I can flip it back for a second, we've come to an agreement with ICANN on performing this study from our consortium partners just recently. I understand there's some confusion about some participants here may be expecting some results from the study, but actually we'll be presenting our study approach to you. We'll be very open and eager to hear your reflections on the approach that we envision, so any interaction would be very welcome during this presentation.

Maybe it's good to have a short introduction of the three consortium partners in this team. To start with, TNO. TNO is the

largest knowledge institute in the Netherlands that is not for profit and organized with the mission to stimulate innovation within the Netherlands and beyond. We're actually with about 3,000 experts in fields that are technical-related, but also in life science research. We are conducting research projects in an international setting for quite a lot. Actually, I think one-third of our turnover is in international projects.

Recently, I've been involved myself also in the scaling the roots study that has being commissioned by ICANN in 2009. Our role in this project will be to contribute our expertise regarding model-based analysis of the DNS behavior, and we will act as a project leader of this consortium.

Maybe, Christian, I can...

UNIDENTIFIED MALE: Sure.

CRISTIAN HESSELMAN: Hi, everyone. My name is Cristian Hesselman. I'm with SIDN. We're the registry for .nl, the Netherlands country code top-level domain. We're currently about 5.6 million domain names, and we have a lot of expertise in managing ccTLDs, also in particular for the security and stability of a ccTLD of that size.

We're also a not-for-profit organization, and the research in this project will be carried out by SIDN Labs, which is the R&D team of SIDN. The folks on that team have quite significant experience in data analysis.

UNIDENTIFIED MALE: [inaudible], NLnet Labs. NLnet Labs is a small, not-for-profit company (three in a row). We do research and software development on Internet technology, and more specifically, we are quite well known for the DNS software we developed and [inaudible] Open DNSSEC.

We also run a number of research projects measuring. We do a lot of measurements with academia on the impact of new technologies on the Internet, that effectiveness of that; barriers for deployment of new technology. And our role in this project, in this context, will be the data analysis, the measurements of DNS infrastructure and data collection.

BART GIJSEN: Thanks so far for the introduction of the consortium partners. Let me proceed with a first sheet on our vision regarding to approach the answers requested in the RFP. Of course, the top two bullets are to provide answers to ICANN as they requested in the RFP. In fact, ICANN requested for a clear view of the current

impact of the New gTLD Program on the security and stability of the DNS, and to identify steps to safeguard the root system security and stability.

In our vision, as you see on the top bullet, we think we shouldn't only do it as a one-time analysis. It should not only be for the current status now, and new gTLDs are being introduced gradually into the root zone, but it should be also forward-looking.

In particular, on the bottom part, we express our idea, our vision, that it should be more or less continuous. In order to be able to create a good playground to get commonly accepted to best practice results, we think it's very important to be open and transparent about this study approach. This is one of the reasons why we're sitting here right now to request any feedback and input from the community.

In addition, a much more elaborate description of our study approach will be published shortly via ICANN channels and be posted for public comments. So if I'm brief here and there in my presentation, rest assured that there'll be more information shortly. Nevertheless, any questions you may have are very welcome here.

We're actually trying to be as open as possible and trying to contribute to a broad consensus on DNS root stability in the context of the New gTLD Program.

Given this vision of our study, here are some of the things that we are actually going to do during the study. Our idea is to be as objective as possible and to make it a data-driven study, as much evidence-based as we can.

In order to achieve this, we will do active measurements. We will initiate DNS traffic and see how the DNS system will respond to that. We will look at log files from, for example, root server operators, but also many other data sources, and we will investigate public data sources such as, for example, within DNS-OARC we [inaudible] the root zone file repository, which is also a source of data that we can use to get some idea on the stability of the DNS and how this changes with the introduction of new gTLDs.

In order to be forward-looking, we thought it's very important not only to stick with data analysis from the current system as is, but we are planning to create a simulation model based on the data analysis, where we will use simulation models in order to extrapolate certain parameters.

For example, if we have seen certain phenomena which are valid for the current system, we may want to have an additional

parameter in the simulation model, which says, “Okay, what if we expand the number of new gTLDs even further? What will be their impact on the stability of the DNS?”

In order to be able to do a continuous study instead of a one-time measurement only, we foresee that we are going to make the instrumentation that we will use during the data analysis and during the measurements, and use that to make them available to the community of ICANN’s community, actually, in order for others also to validate what the measurement results are, and also to be able to conduct further studies for data analysis after we've finished our formal project here.

The conditions under which these instrumentations will be made available will be discussed later in the project. I think it’s a bit too early now to forecast how this will be proceeded to [exactly].

One step deeper into the active measurements, what do we foresee to do there? Actually, one of the instruments that we are going to use there is the RIPE Atlas probes infrastructure. Actually, our consortium has some software framework which allows very easy deployment of new probe-based measurements, and also such that we can repeat measurements over time, over and over again.

Based on this platform, we will be able to do continuous measurements. Measurements will be done from different sets of protocols, which is IPv4, IPv6. Probably also for larger DNS response we'll have to use [inaudible] kind of protocols, DNSSEC. So it'll be validating over a wide range of protocols in order to initiate our active measurements.

In addition to the active measurements, we will use log file analysis and principles, so passive measurements, and we will be collecting data from available sources. So far, we've been having contact with representatives of the root servers I, K, and L. We'll be discussing with these guys on how we can get access, under which conditions we can get access to data.

In addition to that, we'll also use SIDN's .nl data in order to explore a bit on the data sets themselves in order to determine what metrics are valuable and what data we need for that. Then we can play around a bit more with the data sets since they're available within our own consortium.

We're very open, actually, to investigate any other data sets. For example, day-in-the-life data sets or others. If you feel in the circumstances that you can provide us with any data, in particular if it's related to root server operator data, we'd be very interested in chatting with you on what conditions we could get access to those data sets.

Of course, we will be doing some careful planning here on what will be measured and what analysis will be done on these data sets. We'll be defining specific DNS stability metrics. We will also be determining on how much data for how much period of time in order to get some stable results out of this.

I don't know if anybody feels triggered here to provide any data yet, but please feel free to do so.

And then, as I already mentioned, we are looking at public data sources, such as from DNS-OARC, which has stored all the root zone files for I think the last ten years or something like that in their repository. This is accessible for DNS-OARC members, so we will use these kinds of data sources, just as an example, actually.

What we will do is we can extract information from those root zone files, such as what's the root zone file size; how does it vary over time; what's the impact of new gTLDs on this; were any possible errors seen in the root zone files so far, and in particular points in time or for any particular TLDs possibly. That kind of analysis we'll do based on a history of root zone files.

Of course, this is not the only public data source. Also ICANN has its centralized zone data service from which we will be retrieving information to do a similar kind of analysis. And we will be

searching for several other sources, too, in order to get as clear a picture as possible on all aspects of the root zone.

Onto the forward-looking analysis part. What we will be doing is based on the analysis of the collected data shares. We will be trying to find correlations between the data sets. Just to give an example, if we are going to combine active measurements, we'll be inserting DNS traffic into the DNS, and we'll be trying to find those insertions that we did in the passive monitoring data that we'll retrieve from log files, just in order to get indications for correlations, such as what was the end-to-end round trip time of a response, a query response interaction, relative to the number of instantaneous DNS queries that were fired at a particular part of the root zone.

So if you'd be able to find that kind of correlation, we will be trying to find some in variance, based on which we can build predictions of, "Okay, what if the number of queries per second would have increased? What would their impact be on DNS round trip time as seen from outside of the root zone?"

If you have done so, we'll be also investigating several statistical characteristics of those. The ultimate goal of doing this is to be able to do what-if scenarios. What if the root zone would grow in terms of number of gTLDs? But also other kinds of parameters may have a what-if scenario behind them, and this simulation

model would enable us to look a bit further into the future on future possible scenarios.

In addition, based on the data analysis, we are quite sure that we are going to see some uncertainties. All data analyses are always surrounded with some confidence levels and some uncertainty regarding specific parameters that we're monitoring. Not each day will give exactly the same results. The simulation model would also be able to enable us to do some sensitivity analysis of parameters that are simply varying over time or from which the data analysis itself is not quite clear on what the exact numbers are.

So this is where we are going to use our simulation modeling to do what-if scenarios and to do sensitivity analysis on the data.

Now here's a diagram which gives you some indication of the type of metrics that we'll be looking at. The set of metrics is not defined yet, so any further suggestions on where to look for and how to define DNS stability we are very open to. These actually include the parameters that have been specified in RSSAC 002 document. For example, I'll just name a few of those green blocks, which are more or less the DNS stability indicators, something like the publishing latency, the time taken between a notify from the distribution master until that information is available from the root server operator side.

Other indicators I always mentioned, such as DNS round trip time, which you can measure from the RSO perspective, so the time that a DNS query comes in until it's been responded to. But we're going to use our active measurements also to see those DNS round trip times from either some of the recursive resolvers or even beyond that, so more from the client's perspective.

We're also intending to do some DNSSEC chain validations, so we'll try to hook up the DNSSEC crypto-key material responded by the root and validate them towards the new gTLDs that are quite... At least some of the measurements that we will be doing throughout our study and that we will be reporting on later.

Further, once again, additional suggestions for root zone security and stability metrics are very welcome. We're very well aware of the fact that there have been many discussions in all kinds of fora on what the set of DNS health or DNS stability parameters should be.

We'll be looking at a lot of previous studies on this and symposia which they've come up with, and we'll be composing a document on a set of DNS health metrics.

No questions about this? Then I'll proceed with our breakdown of work within the study.

Actually, we separated them in more or less seven work packages. The first work package is including this session, actually. It's involving root server operators in the DNS community. We've planned a number of initiatives to get engagement with the DNS community and to get data collection from root server operators, which then we'll complement with measurements that we'll design and execute, and in work package three, we will analyze those security and stability aspects of it.

Work package four is our simulation model, so the future scenario analysis of the DNS root security. Then we have three more work packages which are more related to the dissemination of the results.

After we have prepared a draft report somewhere next summer, we'll be posting that for public comments, and it will be reviewed by the community as such. We'll be engaging with them in subsequent meetings that will come up either in the ICANN meetings context or DNS-OARC or RIR meetings or whatever. We'll be quite heavily planning some of the feedback moments there. Then we have a project management work package, which wraps it all up, I think.

Given this, we have more or less a time planning, which looks something like this. We're still shifting a bit in where to do what exactly, but at least we've gone off to start by now.

Once again, this work package one is involving the community to get input for the study and feedback on our study approach. Then the work packages two, three, and four are the actual measurements and analysis of data, including the forward-looking analysis with the simulation model. And then with work package five, we'll be preparing a draft report on the study results, and then in work package six, we will be engaging once again with the community on getting feedback and review of our study results.

As you can see, this is taking quite some period, so it's not like we're going to post it for public comments, wait for 40 days, and then get a response. We're actually making an effort to make sure that our study results will be reviewed by the community as such, and we get the opportunity to do some additional analysis if requested.

So this is more or less the high-level time plan for our study approach. That brings me to my final sheet where we wrap up what I've just been saying, I think.

Once again, we envision an objective and data-driven study regarding root server system stability. We'll be looking at the

current state of the root zone stability, as well as extrapolating those results to future scenarios.

We foresee to make our instrumentation available such that we can do continuous root stability monitoring instead of doing a one-time-only data analysis as a result. I guess with our consortium team, which is quite complementary in expertise, we are the guys who can pull this off, actually.

So I don't know if you have anything to add, Cristian or [inaudible]

UNIDENTIFIED MALE: No.

BART GIJSEN: Any reactions, questions, suggestions?

UNIDENTIFIED MALE: I'm [inaudible] from ISP in India. When starting the study, do you have that reference set that how is the performance of the root zone today, so that could become that the reference. Is it becoming better? It may still be good enough, but is it becoming better, versus...

BART GIJSEN:

Right. Thanks. That's a good question. Actually, question number one I think was we're starting as of now, so we don't have any reference data so far, and those questions we did have. We are requesting not only for current data, but also for historic data. So if you're indicating that we need a reference point somewhere in time from which we can see indeed whether it's getting worse or not, that will be on our mind to plot it in time and not to have a one-time-only signature.

We'll be trying to get some historic data in order to be able to even go further back down than the first new gTLDs that have been introduced into root zone already.

How far we can come, I'm not quite sure. It will depend on the data. But we'll make sure that we make an effort to get this reference point.

UNIDENTIFIED MALE:

[inaudible]. So, first of all, congratulations for getting [inaudible] research.

You mentioned you are going to use a couple of different data sources and still figuring out if there are additional data sources. One of the data sources that you mentioned was the SIDN use of the .nl DNS services. Now, I know that the SIDN has a privacy board in place to govern what data is being collected and what it

can be used for. Will your research also have a privacy framework to ensure that the data that you use in your investigation is only used for this purpose?

BART GIJSEN: Cristian, may I pose this to you?

CRISTIAN HESSELMAN: Yeah. We set this up for the .nl data because that's tied to the Netherlands. So we have a privacy framework in place for the Dutch situation, complying with the Dutch Privacy Act. But we will have to look into, let's say, this new situation, which will also involve root server traffic. So this is something that we will need to address within the project, I would say.

UNIDENTIFIED MALE: [inaudible] I was asking...

BART GIJSEN: Actually, I think in our consortium agreement we have a chapter on data sharing policies, actually, which are concerning mainly the SIDN data that we'll be using. I guess when engaging with other parties which will provide data sources, we'll indeed have similar kinds of agreements on which policy applies regarding to data sharing.

Thanks for your question.

UNIDENTIFIED MALE: One of the data sources you mentioned you were going to use is the RIPE Atlas. I think it's a great data source, but it's also a very skewed view of the world. I don't really have a good proposal for an alternate active measurement platform, but I think you should either try to ideally augment the RIPE Atlas network to cover the gaps. That's probably going to be quite different. But otherwise, find some active measurement platform. I know there's commercial ones available, for example.

BART GIJSEN: Thank you, Shane.

UNIDENTIFIED MALE: Indeed. We know the RIPE Atlas well. Well, many of the nodes are in Europe and North America. We do consider other measurement platforms, either the M-lab or the more commercial measurement platforms.. Yeah. We'll take care of for the skewedness of the data, the measurements point.

UNIDENTIFIED MALE: Yeah. Cool. And I have another question. I don't know if it's actually for your team or not. So you explicitly aren't looking at

the provisioning side of things. Is there a separate study underway to look at that? Do you know? Yeah?

BART GIJSEN: Karen, could you respond to that?

KAREN LENTZ: Hi. I'll actually turn that over to David, who's down at the end of the table.

UNIDENTIFIED MALE: A very hot potato, I see.

KAREN LENTZ: David Conrad is our ICANN CTO. Thanks.

DAVID CONRAD: By "provisioning," what do you mean specifically?

UNIDENTIFIED MALE: Well, there was this nice pretty diagram with one part grayed out which said "provisioning." I guess I mean all the ICANN processes around managing this larger data set. It could be quite simple things, like you go to a webpage now, and looking at 10,000 names on a single page or paging through it could be

quite cumbersome. At that point, you start to need better tools, like search, maybe RESTful API that people can do this stuff, and then it becomes... It's not impossible, but there's a lot of additional work involved when things start to scale. That's kind of the area that I was thinking about.

DAVID CONRAD:

Right. So I'm fairly certain — and actually, Karen might be more authoritative here — but I'm fairly certain that a large number of studies on the processes and mechanisms by which the New gTLD Program was implemented are being undertaken. I know there is an internal one done. I'm not sure if there'll be an external one.

Some of that aspect of provisioning I believe has been identified as areas in which improvement is necessary for moving forward. The focus on this particular study is to try to determine whether or not the increase in the root zone has caused damage to the Internet. The incremental improvements in the provisioning system to deal with the additional scale, it might be damaging to ICANN staff or people trying to make use of the glorious ICANN website. However, I don't think it would be qualified in the context of damaging to the Internet as a whole.

UNIDENTIFIED MALE: I think that's probably fair. I have one last thing. I'm sorry. I have the mic, so... You mentioned that you were going to go and look at previous data sets. If you're going to simulation, we have this awesome natural experiment that we've already done, expanding the zone. So is the idea to compare the simulations and use that as a kind of reference to validate them or not?

BART GIJSEN: I think that's an interesting question. I think that thought didn't cross our minds yet. But it's interesting, indeed. Yeah. Thanks.

UNIDENTIFIED MALE: Yeah. I don't know exactly how that would work, but again, you mentioned you have the DITL data. We have the historical. OARC is a great resource for these kinds of things. So that data is available, depending on what kind of simulations you want to do.

BART GIJSEN: Right. Very good.

UNIDENTIFIED MALE: Indeed. It's a great way to calibrate your simulation to validate your model. Thanks.

PETER KOCH:

Peter Koch, DENIC. As such, I am New gTLD-agnostic, but in a previous life, I was part of the team that dealt with what ended up as RSSAC 002. If I remember correctly from your work plan, and maybe I missed something, it goes into 2017. I haven't really seen any interim conclusions there, and given that the data set that was defined in RSSAC 002 was meant to be an early warning system support, maybe there could be earlier results somehow.

Don't get that wrong. Your plans of new and additional measurements is great, and it's very ambitious, and I'm really looking to what you're going to do there. Maybe I just haven't really understood the relation of this early warning system and the missing piece in that, because the data collection is there, but then it's open to the public and needs to be researched. So this early warning system and the study that is scheduled here.

And one final question to add to that is maybe... I've just glanced through the MRFP. These additional sets — I saw there was something mentioned, but that is part of your proposal already that was not really the basic task, right?

BART GIJSEN:

This was not the basic task. Actually, we are starting now, right? So if I paraphrase your question correctly, then you're saying the delivery of study results is quite late, right? That's one of the...

PETER KOCH: I had desperately tried to be more British.

BART GIJSEN: Okay. The timelines you see here are more or less the indicated timelines from the RFP, which includes a draft report or an intermediate report of the study results by the summer of next year somewhere. Actually, those are quite extensive results, I would say, and we will, as I mentioned, make sure that we'll engage with, in particular, the technical community much earlier on, and try to present some of the intermediate results that we have until that point in time.

So, point taken. We will not wait with any presentation of results until summer of 2016, and certainly the 2017 timeline that you're mentioning is in particular in order to allow a long period of, say, reviewing by the community from the moment when we publish our study results.

PETER KOCH: It makes perfect sense to me, at least. So maybe I could rephrase that remark slightly, not to give you a hard time, but to get to the point here. If what we designed back then — and that's, what, already two or three years ago that the work started — f that it meant to be an early warning system, maybe you find some

indicators whether it is really able to serve that purpose. That should be something that might be in your earlier results.

Given the ambitious project that you've mentioned, I'm fully aware that that will take some more time, the measurements and all the results. Thank you.

BART GIJSEN: Thanks. Any further comments on this?

UNIDENTIFIED MALE: No, just in addition to that, indeed, we will do outreach and present intermediate results insights during DNS-OARC meetings, ICANN meetings, RIR meetings. It's part of our project plan, and we'll try to involve community as much as possible and get feedback.

UNIDENTIFIED MALE: I have one comment. [inaudible], ICANN .cc. I'm in charge of K-root and ATLAS as well, So that's somehow related. Basically, it's a little bit of nitpicking, but when I come to ICANN meetings, they tell me you have to be a nitpicker.

The name of the project is Analysis of Root Zone Stability, which for me as a root operator basically doesn't have any real meaning. I saw the presentation, and basically your take on that,

which I agree with, was root server system stability, which is fine, but root zone stability, yeah, at best can be like if the file is published properly.

So I think it's good to start with a proper name from the very beginning. It's not your thing. It's ICANN's job of naming the project. But I think it's good to change that.

As a second thing, for those reviews and all that stuff, I strongly suggest that you also involve the RSSAC Caucus, because we have a group of almost 70 DNS experts in the caucus. Most of you are members of that caucus. These are experts who can give you real, valuable feedback.

DANNY MCPHERSON:

I just had one question. I guess from what I understand this is focusing on the stability of the root server system itself and the implications of, say, either [inaudible] or additions to the root zone file, and not the potential effects on Internet security or stability of relying parties in the root server system. For example, most of the things that were hung up with name collisions, internal name certificates, public suffix lists, other things that resulted in some of the challenges with deployment of new gTLDs. Is that correct that that's entirely out of scope and we're just looking at the root server system and the root zone file?

BART GIJSEN: Yes, indeed. We don't look at relying parties in this study. Any comments or suggestions are welcome, of course, but it's mainly focused on these two systems, yeah.

JORDAN BUCHANAN: Hi. Jordan Buchanan with Google. I'm more confused after having seen your previous... Can you go back to the chart on all the things that are in scope and out of scope? That has recursive resolvers on it, which are not part of the root system. So are they in scope or out of scope?

BART GIJSEN: I guess they're not in the primary scope. Nevertheless, we'll do active measurements, which are from the client side of the recursive resolvers. So we'll be seeing some of the behavior of the recursive resolvers.

Having said so, once again, they are not in primary scope. They're not the subject of study, but they're that closely related to the root server system that we cannot ignore it. Okay?

JORDAN BUCHANAN: Sure. I guess I just imagine things like cache fragmentation are going to have an effect on the recursive universe as well, and it

would seem like those are just as important to the DNS scalability topic as the root zone servers themselves, but maybe that's more a comment for ICANN as opposed to for you guys.

DAVID CONRAD:

Yeah. Actually, the office of the CTO will be working in conjunction with TNO on this project in the areas such as the implications of how the scaling of the root impacts caching service]. And that is a topic, while it's not directly in scope for TNO, it is in scope for the stuff that we do within the offices of the CTO, so that's one of the research areas that we're going to be looking at, and all of that will sort of be put together by staff at some point when we're sort of aggregating the information associated with the studies that were essentially commissioned as a result of the NGCP directive or whatever, however all this stuff started. All of that will be put together at that point.

JORDAN BUCHANAN:

I have one more question or more like a comment, I guess, which is to build on the previous gentleman's comment about the early warning aspect of this, which is, in addition to the RSSAC 002 I think the Board has passed a resolution that basically says we're not going to allocate any more gTLDs beyond this set until we see the results of basically what our current...look at our current monitoring situation. And it seems

to be phrased of the context like “early warning” as opposed to necessarily the full root server studies.

So it seems like there’s some direction from the Board that we want to see some early measurement of the issue as well to make sure that we’re not doing active damage before that process proceeds as well. So it would helpful from that perspective, too, I think.

BART GIJSEN:

Thanks. Any further questions? If not, then, Karen, I think we can conclude. Thank you very much, and we’ll be looking forward to interacting with you on this interesting topic. Thanks.

[END OF TRANSCRIPTION]