
HYDERABAD – Root Stability Study Workshop
Tuesday, November 08, 2016 – 14:00 to 15:00 IST
ICANN57 | Hyderabad, India

UNIDENTIFIED FEMALE: ICANN 57, November 8, 2016, 14:00 to 15:00, Root Stability Study Workshop. ICANN 57, November 8, 2016, 14:00 to 15:00, Hall 6 Root Stability Study Workshop open session.

ELEEZA AGOPIAN: I think we're ready to start. Are we recording? Okay, great.

If anyone who is sitting in the audience would like to join us at the table, you are more than welcome. We'd like for there to be a discussion once the presentation is over, so please feel free.

I'm the Manager of Operations and Policy Research at ICANN. I wanted to tell you a little bit about the Root Stability Study (or CDAR) that is going to be presented here today, which was published for public comment about a week ago, on October 27th.

Could you advance a couple of slides?

To give you a little bit of background on the study, it was commissioned following a couple of recommendations that were made to ICANN and that was adopted by the Board, asking for a

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.

review of what the new gTLD program's impact was on the security and stability of the root zone system.

The board committed to conducting this study and committed to postponing any further rounds of new gTLDs until there was clear evidence that there had not been any degradation in the system. In 2015, we issued a request for proposals and contracted with TNO, who's joined up here on my right. Bart Gijsen from TNO is here, who's going to be speaking with you about the study.

TNO is working in a consortium together with SIDN and NLnet Labs. Cristian and Jaap from those two organizations are here with us today, too, and they're going to go through their findings. I think I'll turn it over now to Bart.

BART GIJSEN:

Thank you, Eleeza. (I'll use this pointer here.) My name is Bart Gijsen, I work for TNO, which is a knowledge institute in The Netherlands. I'll be walking you through the draft report that you can find for public comment. It's been published last week.

I'm here together with Cristian Hesselman from SIDN Labs (the research organization from the registry for .nl) and Jaap Akkerhuis who is one of the experts who was already involved in previous scaling the root studies, just as I have been myself.

As I said, I will be walking you through the highlights of the report. Any more details on the precise study, you can find in this report.

So what was the context of the CDAR study? The diagram that you see here on the screen demonstrates how the root zone size, in terms of the number of TLDs which were included, has expanded quite a lot when delegations of new gTLDs started in October, 2013.

The diagram here illustrates that this is kind of a disruptive trend. This may have caused some questions or concern about what could this do to the stability and security of the root DNS system. This is exactly the research question that we've been investigating in this study. Did the delegation of the new gTLDs degrade the stability or security of the root DNS system?

Using results from that data analysis that I'll be walking you through in a minute, the secondary question is, are these analyses giving any rise for expectations that there might be a degradation in case the root zone would be increased even further?

So, what did we do? The analyses that we did are largely based on a number of data sets which contain large amounts of data regarding DNS queries towards the root zone, but also active measurements on requests and the response from the root DNS

system, primarily using, for example, the RSSAC 002 measurements which are being collected and aggregated by root server operators.

They have been collected since October, 2013, so more or less around the initial delegations, the first delegations of new gTLDs. We also used DITL measurements (or Day in the Life data collection events) that have been going on since 2006. And if I recall correctly, we've, in particular, used those data collections from 2012 – so prior to the initial delegations of new gTLDs until recently.

And these Day in the Life measurements are two-day data collections of all queries being sent to root server operators, and they share in this large data collection exercise and, via DNS-OARC, this kind of data is made public or at least made available to DNS-OARC members.

The RSSAC measurements and the DITL measurements are like measurements from inside the root DNS system itself, so it's data collections by root server operators.

When measuring from the outside, we used RIPE Atlas Monitoring Framework. RIPE Atlas is probably the largest active probing mechanism which continuously sends DNS queries to root server letters and receives responses. Based on that, we can

retrieve an insight on the behavior of the root server system from the outside on.

These are, I think, the three core publicly available datasets that we used in our analysis, but we used many more other sources, such as the zone file repository, ICANN's gTLD registry report, which contains statistics about domain names per TLD.

We used incidental datasets regarding the renumbering of h.root, and we also used some of our own tools – in particular an NLnet Labs tool, the DNSSEC validator – in order to check whether DNSSEC validation errors would occur for new TLDs, if any.

Based on these datasets, we investigated the stability and security of the root DNS system, and then we had to drill down a bit in what kind of metrics do you measure if you're talking about the security and stability of the DNS system.

Quite frankly, there's not a real consensus on what exactly we understand under the security and stability of the root DNS system, but there's a bulk of work being done by the technical community to try to find definitions of the root stability.

Here's more or less a breakdown that we ended up with. We regard the stability and the security of the root DNS as being part of two distinct types of metrics: the operational stability of the

root DNS system; operational stability including, amongst others, the query rates of the load which has to be handled by the root DNS system, as well as how fast or how – can it respond under all circumstances to the queries being sent to the root DNS system. So then, we are more talking about the query and response stability of the root DNS system.

I'll be giving some more details on this in the next few sheets. Apart from the operational stability – so the ability of the root DNS system to respond to the queries that are being fired to the root – there's also the question of, is the data that's being returned consistent?

And of course, we experienced this as being the case, but we wanted to do some checks anyway and to see whether new gTLDs would have a negative impact on the data consistency. For example, would other kinds of errors occur, or data inconsistencies occur in the root zone file; or new gTLDs being signed via DNSSEC. Do we see more DNSSEC validation errors – yes or no – for new TLDs than for older TLDs?

These are more or less the metrics that we investigated. Just to jump to our primary conclusion here, investigating those large data sets didn't show any degradation of the stability or security metrics that I just presented to you that can be attributed to the delegation of new TLDs.

In the next few sheets, I'll be giving you some more findings on which this conclusion is based, but this is more or less our final finding.

Finding number one: the traffic towards the DNS root system is continuously increasing. What you see here in this diagram is a plot from RSSAC 002 measurements which show you the query rate to individual root server letters.

These are shown in the separate graphs here in the separate colors, and they show you that, on average, the query rate towards the root DNS system is increasing. Having said so, an increase is not very [monotonous]. It's not a gradual increase over the course of time. There are fluctuations around the query rates.

There's not a thing as one single definition of what exactly the query rate is. For some root server letters, the DNS query rate might be increasing over some periods of time; while for others, it's slightly decreasing. But in general, we would say the traffic towards the root DNS system is increasing.

Along with this increase in traffic, as you may have heard in the public meetings here from RSSAC, the amount of servers which are included in the root DNS system are also growing. Apparently, given our conclusion, the root DNS system has been

able to grow along with the increase in root server traffic quite well.

Having said this, what's the relation to the new gTLDs? There, we have to drill down the total query volume in a number of components. Actually, we distinguish between normal days in the internet and, say, the abnormal days in the Internet.

The abnormal days in the Internet we experienced, for example, on November 30 last year and June 25th of this year where we've seen an unusual high number of traffic towards root server letters. And as you may conclude yourself from reports that have been written about the causes of these events, there's no relation whatsoever to the introduction of new gTLDs. So, I'll leave those unusual incidental peaks of query volumes towards the root DNS system out of scope here for the time being.

This leaves us with the regular days in the Internet. There, we can distinguish between queries for valid TLDs and the nonexisting, so the invalid TLD queries.

Simply put, valid queries are those queries that are for TLDs which are in the root zone file at the time they're being queried, and invalid ones are all of the queries ending up at the root DNS system.

Now here's a decomposition of what the valid and invalid query fractions are towards the root. What you see in this diagram is five more charts showing analysis of the number of queries to valid and invalid TLDs for the DITL measurements in 2012 until the DITL measurements in 2016.

The top three – so the purple, the red, and the green bars – in this bar chart are for valid queries. The purple is for all the queries that are being sent to .com. As you see, in 2012, this was about a little bit more than 20% of all queries to the root DNS system; whereas you can also see the delegations to other TLDs.

In green here – and that's for those who have really good eyes – you see the fraction of the number of queries that have been querying for new gTLDs. Those are being introduced at the end, so after the DITL measurement in 2013, so you can only see them in the DITL measurements for 2014, 2015, and 2016.

And as you see – or actually don't see – the fraction of queries for new TLDs is very little. In particular, in the DITL 2016 data collection, we view a 1.1% of the total valid query rate – so of the top three colors in this diagram. 1.1% of the total valid query rate is for new TLDs; whereas this only amounts to 0.4% of the total queries that end up at the root DNS system. So, traffic directly related to the new gTLDs is very low.

It is slightly increasing with the increasing number of TLDs in the root zone file, but it is rather insignificant so far.

So, didn't we see then any effect of the new gTLDs? Not exactly. We took out our microscope and dove into the data once again, and looked in particular at points where new gTLDs were being delegated towards the root zone file.

The four plots that you see here are time based plots in number of weeks, so it's about a nine weeks' period of time for each of those plots.

And on the left-hand side, on the Y axis, you see the number of queries ending up – for example, a particular new gTLD.

What you see here in the top left, the diagram is a rather typical behavior. The black arrow indicates the day at which this new particular new gTLD was delegated to the root zone.

So, prior to and to the left of this black arrow, it was not yet delegated in the root zone file; however, it was queried for. So, queries are being sent to new gTLDs even prior to their delegation. In fact, the root DNS system receives more queries for new gTLDs before they're delegated to the root zone than after.

I'm not pretending to have the exact answer here as to why that happens, but it's highly likely that after delegation, the resolvers

kick in. In particular, the queries are being cached in the resolvers, so it's that for next queries being sent for new gTLDs, they will not end up at the root DNS system anymore. They will simply be responded by the resolvers. So, the resolvers help here to decrease the number of queries for those kind of new gTLDs.

On the right top side, you see more or less the same kind of picture. You will also see in these pictures that the query rates converge rather quickly to the new situation after delegation.

As you can see in the other plots, there are some other typical behaviors too for particular new TLDs. But once again, look on the left-hand side regarding the query rates that are being sent to those new TLDs. Those are, on the root DNS system scale, very small. So even the large peak that you see on the right bottom – and you might think that that's a huge spike, which it is in the relative sense. But the number of queries that are associated with the time of delegation and the peak that you see there – the number of queries is, for the root server system, very low.

To investigate what these kinds of microscopic behaviors do with the experience of the root server system behavior from the outside, we used RIPE Atlas measurements here. For example, we looked at the response times during those periods of delegations.

Recall that we saw some peaks, and we might guess that for those particular new gTLDs, the response times might increase because there's a slightly higher load. These diagrams here illustrate that actually, the period of delegation – the actual time period – is exactly in the middle [inaudible] of those diagrams.

Those diagrams show you an hourly plot of 100 hours' duration, querying the response times for several root server letters. And exactly in the middle of those plots, a delegation of a new gTLD happened. So, what you see is that it actually doesn't even ripple the response times of the root DNS system.

So, if you look from the outside, even during those periods of initial delegations of new gTLDs, there's nothing to be seen, actually, in terms of those times, which means that Internet users – or, in particular, resolvers – wouldn't really be able to even notice that new gTLD delegations into the root DNS system have taken place.

I promised you a word on DNS data consistency. We scanned for several types of data inconsistencies, but we were not able to find any. What we, amongst others, did is we checked whether different errors or more errors would be introduced in the root zone files, actually, because maybe some of the errors might slip through.

I have to say here that the whole process of producing a zone file, distributing it and publishing it by the root server operators is a very well thought-through exercise, and it's meant to make this whole a zero-error system.

Now, having many more new TLDs, could that change anything in the error behavior of the errors in the root zone file? Actually, no. We've scanned a lot of them. We've queried some of the people involved in this whole process of producing the root zone file, and there's no additional errors that are being introduced by the new gTLDs and delegations.

Similarly, all of the new gTLDs are being signed via DNSSEC. Could this lead to more validation errors, for example? Could this lead to broken chain events that we have not seen so far? We've done some checks on this, actually. We have a monitor running on verifying whether the [parent] key change between the root and TLDs – actually all TLDs, but also the new TLDs – whether there are any validation errors that we detected there.

And in some cases, when new gTLDs were first delegated to the root, we saw some errors. But those were quickly resolved and there are no more validation errors seen from new gTLDs, although there are many more new TLDs than there are other TLDs. So, we didn't conclude any issues here regarding DNSSEC validation.

Those are the findings based on which we conclude that we didn't observe any degradation of the stability or the security of the root DNS system that are due to new gTLDs. Of course, we recognize that this conclusion is confined by some of the imperfections that we have in publicly available measurement data.

It is also confined by the particular choices that we made and the design that we did of the analysis that we executed. But having said so, we also had many interactions with the technical community regarding these findings and even preliminary findings, and so far, nobody stepped up to us and said, "Okay, you forgot about looking at this particular effect which might have a negative effect on the introduction of new TLDs on the stability of the root DNS system."

But once again, we are very interested in hearing whether any other concerns are still living out there.

Then the other question would be on if we take a look at the data and observe what we've seen in the data sets so far, what does that mean in terms of what can we expect for potential future impact if the number of gTLD delegations would be expanded even further?

The first remark to be made, which we can base on the data that we saw so far, is that if we presume that the evolution of new

gTLD delegations would continue in the current pattern, and if we further presume that we have some time invariant correlations that we see in the data currently – and those have been time invariant for the past years – if we assume that they remain time invariant for the future, then no. We see no signs that adding more new gTLD delegations will degrade stability or security of the root DNS system in the near future.

Of course, these depend on these presumptions that we make, so what we also did is some more speculative analysis of what are the additional risk factors that these presumptions are gained invalid. We did some brainstorming on this. We didn't put too much effort in it because it is rather speculative and we wanted to focus on our data-driven approach here. But nevertheless, we did some brainstorming on what possible impact that could be in the context of future new gTLDs delegations.

One of the things that could be argued is that a risk would be if a handful of new gTLDs would grow very fast in popularity towards .com-like popularity. That would render some of our results as invalid, I guess. The likelihood of this happening? I don't know. .com also took several decades to grow to its size. At this point in time, we don't see new gTLDs growing that fast. But once again, if this would happen, that might render some of our presumptions as invalid.

Similar [holes] if the growth of the number of new gTLDs would be unbounded. So far, we've seen introduction of hundreds – or say 1100 – of new gTLDs. If this would have been in the thousands of new gTLDs, I don't know. I'm not quite sure whether we would see exactly the same results as we've seen now in the data.

Similarly, there is a potential risk at this point in time. As I've shown you, the number of invalid queries to the root is currently exceeding the number of valid queries to the root. One of those reasons is that .belkin, .home alike queries which are really not cached because they're not valid TLDs.

If we would have a new gTLD that would grow very popular and would then, for whatever reason, need to be removed, in that case – it's a rather hypothetical case I guess – we would have yet another .home-alike effect, which could create invalid queries to the root DNS system.

Once again, all of these, in our opinion, are rather hypothetical exercises, but are risk factors anyway. We couldn't come up with more concrete risk factors that have a higher likelihood, but once again, we're open to hearing any other suggestions that you might have.

And then there are a number of risk factors which are in our opinion actually not new gTLD-related, which might have a much

bigger impact than the introduction and delegation of new gTLDs, and that is if we see protocol changes, that might increase the amount of processing on the root name servers.

For example, if we see a shift in traffic from UDP to much more queries TCP-oriented, that might really increase the processing power on the root DNS system. If that happens very fast, then it might be a risk to the scalability factor of the root DNS system as a whole.

However, as I mentioned, those are not new gTLD-related, and once again, the likelihood of that happening, as far as we are concerned – and it's rather speculative, but it's not very likely.

Having said so, these are our final recommendations on the further expansion. We would recommend to remain the enforcement of the current measures that have been taken in the new gTLDs program and keep them as is.

Once again, if we are capable of having an evolutionary development of new gTLDs and delegations of new gTLDs, then we are quite sure that the stability and security of the root DNS system is not a danger in any way.

So, measures that have been taken, if you continue those, that will help in terms of keeping the evolutionary path instead of having a disruptive trend somewhere. In particular, these

include, for example, the gradual rate of delegating new gTLDs. If you continue that, then you have one way of reassuring that the current pattern will simply endure to exist.

Also, I think it's very wise to keep monitoring the impact of new gTLD delegations. But in general, actually the more wider risk factors that we identified, it's good to have a more continuous measuring mechanism for observing the impact of new gTLDs and the other risk factors.

That's, I think, more or less what we have to bring to the table here. Just before I go into any questioning and to make sure that you understand where we are in the study, we're here to collect any of your feedback. You can also provide us feedback in terms of the public comment period, which ends by December 22nd. We'll use this feedback to produce a final report planned for April of 2017.

Finally, I'd like to recognize some of the providers of relevant datasets that we've been gratefully using for this study in particular. DNS-OARC, RIPE NCC, amongst others. And also, those persons how have made it possible for us to interact with the technical DNS community, including RSSAC, SSAC and many other institutions in the DNS technical community.

This, more or less, is a thing that I'd like to mention to you, and, once again, we are very open for any questions.

ELEEZA AGOPIAN: Thank you, Bart. I'd like to open the mic. Rubens, go ahead.

RUBENS KUHL: I'm with .br. I'd like to start from the point of the gradual delegation of TLDs – that even though it's sound advice, it seems more like a common sense advise because there is no basis in the report for recommending that. So even though it's something that I agree, it's something that has no counterpart in justifying it.

One of the problems in that is that because there is no research to back that up, we also don't have any guidance on what levels would be used or what rates should be used. Once upon a time, one thought, "Hey, let's think of an imaginable number of TLDs, like 1000, and say, 'Hey, let's do maximum 1000 a year.'"

That had no basis whatsoever, but some might say, "Hey, that's any imaginable number," and that end up being prescribed as, "No, that's the limit in the root system. That's the capability of the root system to absorb." And that's not true.

So, prescribing a gradual rate but not recommending one possible rate or conditions to rates, or possibly stepping systems like, "Hey, do 50 and then if nothing goes wrong, do a run," and

not prescribing any of those is something that I found strange in the report.

I also note that you end up suggesting to not retire TLDs. What would be needed to implement that? Delegate every retired TLD to ICANN that could run a NIC.tld website and say, "Hey, this was retired. There is no more such TLD?" Is that what the prescription is for? I had some other question, but I'll let you circle around for other people questioning, then I'll be back. Thank you.

BART GIJSEN:

Okay, thank you, Ruben. Interesting questions. To go into the first question that you have, why the recommendation of a gradual increase of delegation of new gTLDs that would have no foundation in the report, and in particular, also, that the quantifications that have been made previously on the 1000 a year, that those are not based on any research.

I partly agree, and I think what we've shown in our report is that if the current way of evolving the root zone with new gTLD – if that is sustained, then we have found evidence that there is no real reason to be worried about the stability and security of the root.

So, any kind of measure that has been taken, if you continue that towards the future, will contribute to remaining this evolutionary

path. And once again, if there's this evolutionary path, then we see data which says you will not run into any stability or security problem. So, that's where this recommendation comes from.

Now on the quantitative part of this, we wisely did not use any rates of delegation. We merely said the measures that have been taken, and I think 1000 a year – you can debate on whether that's been argued by any research. But I think at this point in time, the rate of delegation is simply limited by administrative processes.

As long as we keep that in the same rate, that contributes to any of the evolutionary path. Do you want to kick in here, Jaap?

JAAP AKKERHUIS:

Yes. Basically, what we say there is that we should err on the conservative side. Talking about rates, I think the first delegation started about a year ago, something like that, but the real bulk of delegation started pretty recently, and we actually have a rate of 1000 per year, [roughly], doing it by hand.

And when there's no data, there's nothing to predict. So, erring on the conservative side is actually the pattern we see.

BART GIJSEN:

The other question you had is on not to retire TLDs. That's actually not one of our recommendations, right? The only thing

that we said is there could be a risk if you would get a very popular new gTLD, and for whatever reason it would need to be retired, then you may want to be sure that you don't run into a .home or .belkin kind of situation.

We don't make any recommendation on not retiring any new gTLDs or something like that. Right? Any other question here?

UNIDENTIFIED MALE: Let's circle around and see if more people have other questions.

JOHN LAPRISE: I want to thank the group for doing some really fine research. I just have one question. The report notes that it's unable to identify the causes for a number of phenomena, including the growth in the number of queries that are sent to the root. Is there some sort of incentive or some sort of looking forward recommendation to examine the causes, just because we sort of want to know what's going on?

BART GIJSEN: Right. I think from the tech community, we're all in it to understand what we see, right? We were actually quite happy to get the opportunity to have a look around some more thoroughly; but, given the opportunity, we also saw that there's

huge amounts of data and huge amounts of effects – a lot of effects going on.

The Internet isn't stable, for example. It varies day by day on the technical side, on the querying side. There are events going on, like the 30 November event. So, there's that kind of tremendous number of things that are going on on the internet, which make it very hard to explain in detail what we exactly see.

Now, I think – I hope – in the report, that we made some findings which build some additional expertise and insight that we have in the way it works and why it works that way. But we wanted to be relatively cautious and we want to make sure that this report will not explain the Internet and the behavior that we see on it. So, that's what we'd like to emphasize here.

I don't know if any of the other consortium partners has any additional rebukes to this. Not really.

WES HARDAKER:

I'm with USC/ISI, b-Root. The rate has been increasing since before the new gTLD program. I don't have a graph to show you and I haven't compared the rates before and after, but it's not necessarily just because of the new gTLD program. So, a much longer period study would need to be done to really draw a conclusion.

BART GIJSEN: Thanks, Wes. I agree to that. During the study, it's time [inaudible] based, and the Internet, indeed, is continuing to grow and the root zone is continuing to grow. This is where our remark comes from to keep monitoring, and hopefully a bit more continuously monitoring of the impact of new gTLDs, but also other effects. Any other questions? Peter?

PETER KOCH: I'm with DENIC. Maybe this was out of scope for the report of the study, so I apologize in advance, but when you relate the queries received for new gTLDs versus all the other TLDs and you see that there's only a small contribution, would that allow for any assessment of the necessity and/or success of this controlled interruption thing, as in name collisions?

JAAP AKKERHUIS: Well, it's the last words you say. You say "name collisions." If you see that the non-delegated domain which has been suspended for by the Board to be delegated for a long time – I don't know of exactly the terms. The big one is .home, and it already has 5% of the total rate. It's clear that it is way more than all the 1500 together.

So, there is – at the moment this [is going to be launched] as gTLD there will be some problem. That's more to do with cause of the name collision there, if you get my drift.

BART GIJSEN: I must say, this is a very bad spot to stand to receive the response. No, it was not your fault. I should maybe walk five meters back or so. But I sense the response was kind of –

JAAP AKKERHUIS: Well, it kind of depends, you know. If you now [would release] .home, the name collision will really kick in and [been interrupted.] Maybe this interrupt is saying, ["I also get the name."] You will see a lot of things there, but for a lot of them – I don't know; it's difficult to say because the load is so small.

PETER KOCH: Yes, that was part of the question. I wouldn't start with a big elephant there, of course. If I remember correctly, there was a list that was sorted and then there were some people – say arbitrary threshold, or whatever, rule of thumb threshold – drawn for the further treatment. And again, you might be the wrong panel to ask this, but did you look at any of these in particular?

JAAP AKKERHUIS: No, we didn't really look at – we only noticed that some of the domains mentioned in the name collision report had to indeed be [inaudible] That's why we kept them separately. One of them is .home.

PETER KOCH: Thank you.

UNIDENTIFIED MALE: Good afternoon, everybody. [I have a question.] I come from China and our company [softmeds].corp [inaudible].home. I'm also asking a question. I think this is a technical problem for the name collision, so I think that with a high-speed technology development, this can be handled, and [well controlled].

So [inaudible] is not an excuse for [throwing the three] new G[TLD] further in the air. I think we can look forward and make some more study to evaluate the safety. So, I need more quickly reply for the [three] new gTLD. I'm looking forward it can be reopened. Thank you.

BART GIJSEN: Thanks for your response. Actually, we are the wrong team to lay this question at. It will be noted here, and I guess it's in ICANN's hands to handle this, right?

JAAP AKKERHUIS: Although [inaudible] also has quite a big query load compared to all the new TLDs together.

UNIDENTIFIED MALE: I mean, with our group can study your paper to declare [this is safe] [inaudible] it will not [inaudible] for all the – how do you say – it's not difficult to open the two – the three gTLDs. I think we can do much more work on study on the name creation problem at the basic rooting system, the DNS rooting system.

BART GIJSEN: Right. But this point in time, we didn't have any focus on particular new gTLDs – right? – and looked more into the wider aspect of the impact of delegated new gTLDs. So, I think you're raising relevant questions here, but they're slightly outside of the scope of our study. Right? So, if you would have any particular questions regarding our report, I'd be more than happy to answer those.

UNIDENTIFIED MALE: Okay. No, thank you.

ELEEZA AGOPIAN: We have a remote question.

UNIDENTIFIED FEMALE: The question in the chat is, "The new gTLD delegations increased the invalid queries (i.e. NX domain attacks) and also invalid subdomain queries attacks of valid domain on new TLDs. So, what security measures have been taken to avoid these?"

BART GIJSEN: Let me try to read this one again. Actually, the question starts with "new gTLD delegations increased invalid queries...non-existing domain attacks." Actually, this is not really what we have observed. Let me try to answer this one.

The only thing that we see is that prior to delegations, of course, there are queries which will be responded with nonexisting domains. But as I've shown you in the diagrams there, even those query rates are very low if they are directed towards new gTLDs. In fact, those are a fractional portion, I think, than also for new gTLD queries.

UNIDENTIFIED MALE: And also, expanding on Wes' feedback, I think that the amount of invalid queries in the root has been growing for a number of years as well, so I don't think that we found that it could be attributed to the new gTLD program.

WARREN KUMARI: Kind of responding to the previous question – yes, the delegation of a new TLD in no way seems to increase the number of queries to the root. In fact, if anything, it's likely the opposite because it removes one string that could be queried for.

But the second part of their question was, is there anything being done about stuff like random subdomain attacks or NX domain attacks on TLDs, I guess. I missed part of the point of the question. I mean, there are things like aggressive NSEC and things like that which should significantly cut down the amount of junk queries.

If everybody did it, it would drop from like 65% of queries to the root to less than 1%, and the same thing happens in the TLD zone. I'm guessing that's what the question was about.

JAAP AKKERHUIS: One part of the question is whether there should be administrative restrictions on a meter-based restriction on the flow rate of new TLDs, which sounds to me while you have kind of limit and the moment you hit a limit you don't do it anymore.

That's not what we're saying. We're saying that currently, we don't see a problem. But since we don't see a problem, we also cannot say when the problem would be. So, that's why we

recommend to measure what it is. So, you find time when some things are getting too fast and getting out of hand. That's actually the basis of what we say.

And yes, we hope for [inaudible] things like aggressive negative caches, stuff like that will actually decrease the invalid response. Another theory – but that's just speculation – is that since there are more and more big public available resolvers appearing on the market and being used – things like the Google [Name Cache] and [8.8.8.8.] – is that they will actually catch more of the legal queries.

That's why you see more of the illegal queries popping up at the root, since the negative caching is always sorted and [forward caching.] You really have to dive into that part of the [inaudible] before you can say anything there. But maybe Warren wants to answer something about that.

WARREN KUMARI: No, Warren was just going to say that it's really hard to hear in this room. The audio seems weird, like I heard four words.

JAAP AKKERHUIS: Do you want me to try again?

RUBENS KUHL: You mentioned during your presentation of one string that was applied for in the last round, and you mentioned that it would probably deserve treatment similar to .home, which was .belkin. Doing your study, did you find any other strings with that behavior, besides .belkin that you already mentioned?

JAAP AKKERHUIS: Well, we didn't really look to that, but we found some interesting strings. .corp is one of them. But this is not part of the study, again, and if you really want to know about that, it requires some new analysis of the data there. But then we're going again into more the name collision part of the [inaudible] than just new gTLDs.

RUBENS KUHL: You also mentioned that after delegation of TLDs, it converged quickly to a steady state. Was that like in hours, weeks, days?

BART GIJSEN: As you see here in this diagram, I think where you're referring to, mostly it's the day after delegation, actually, and the actual data sometimes takes slightly longer as you see on the left bottom diagram. But particularly, the system converges very fast, within a day.

RUBENS KUHL: Okay, thank you.

BART GIJSEN: Other questions? No. Eleeza, would you like to wrap up?

ELEEZA AGOPIAN: If there are no other questions, I think we can end the session.
Thank you very much to the three of you for being here, and
thanks all for coming.

[END OF TRANSCRIPTION]