

---

PRAGUE – GAC / SSAC Joint Session  
Tuesday, June 26, 2012 – 14:00 to 14:45  
ICANN - Prague, Czech Republic

PATRIK FALTSTROM: SSAC members that are in the room that think they have something to say, which I hope is about everyone, please be seated at the table where there is empty space.

CHAIR DRYDEN: Good afternoon and welcome back everyone. I hope you had a good lunch. As you know, we are now meeting with the security and stability advisory committee. So Patrik Faltstrom, to my right, is the chair of the SSAC. And, as usual, we are appreciative of you taking the time to come and meet with us. We have a few suggested agenda items to discuss today. So I will hand over to Patrik, if I may.

PATRIK FALTSTROM: Thank you very much. I have with me to my right Jim Galvin, vice chair of SSAC and also chair of the membership committee. I also have several GAC members in the room. Of course, I have. There are several SSAC members in the room.

[Laughter]

And I hope that -- in the discussion that we had before this meeting, we had an agreement to (audio problem) -- these discussions.

---

*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.*

---

We talked. And it was the chair, the vice chair, so GAC and Stefano, together with a few individuals from SSAC that looked at the various issues that could be interesting to discuss. And we came up with what you see on the screen, those three main topics.

So, if we start by informing you about the law enforcement/SSAC collaboration that we are doing, we had quite a long meeting with law enforcement yesterday, Monday of the ICANN week, where we discussed specifically two topics. One had to do with validation of WHOIS data, which is validation of data that is registered and tied to domain names.

The second thing we talked about had to do with carrier-grade NAT and various address translation mechanisms when they are in use by Internet service providers.

In both of these cases, law enforcement explained what the problems were from their perspective with the lack of, if I may use that term, lack of proper validation today and also the problem with tracing where various attacks or transactions over the Internet is taking place, if it is the case that there is -- that there are network translation boxes in the network, which means that it's hard to know who a certain IP address belongs to.

The address translation implies that the address -- IP address in use is actually changing in this box. And that means that it's hard to know what the actual -- who -- it's hard to request data for the individual at the other end.

---

We created this agenda in cooperation between law enforcement and SSAC by creating a small group of people just like we did with the agenda for the GAC meeting. And we decided for the meeting in Toronto we're going to work this in just the same way, that we are going to put a small group together that talk about the various discussion items to talk about.

The second item -- I was thinking of going through all three and then open up for discussion.

The second item had to do with batching and digital archery where the question, of course, from -- where there was some questions from GAC regarding the need for batching, the risk that exists on the root server system, which are related to root scaling and the various scaling reports that are done that we have been talking about before. We in SSAC, of course, have been talking about this during the last couple of weeks and, specifically, the last couple days since digital archery -- since it was decided that digital archery was not to -- was to be suspended.

We want to clarify that, from an SSAC point of view, it is not the case that the number thousand that is talked about is something that is a hard number. Basically, it is wrong to say that 999 TLDs can be introduced in one year and that works and 1,001 will not work. Instead, as is written in many of those scaling reports, it has to do with the rate of change and how the root server system as a whole can absorb the new TLDs, which is the important issue.

We do believe that it's a good thing that we already have deployed DNSSEC and IPv6 in the world, because there were some unknown -- there were some unknown risks with -- adding new TLDs at the same

---

time as -- new TLDs at the same time as IPv6 and DNSSEC. One can call it the communitary effects. And those are, to a large degree, gone, which is a good thing. It's still the case that we feel that what is needed is some kind of feedback mechanism so that, if it is the case that the service that is given from the root server system, for example, when someone wants to change a name server record, an existing TLD want to change a name server record, if it is the case that the time it takes to get that change implemented, just because the request happens to come during the period where new TLDs are introduced, then we believe that the introduction of TLDs should slow down. So it's that kind of sort of feedback loop. And, because of that, with the help of a feedback loop, it might be the case that it's possible to add new TLDs faster. Or it might be the case that you have to do it more slowly.

So the feedback loop is the important thing here. Because there are so many involved parties that all of them has promised that they will be able to handle this and that. And, of course, we don't see -- there's not anyone that do not believe in the various promises and what people are saying they're going to do. That is absolutely true that it will work. But we cannot risk that something is happening. And, if we see the earliest - - an early indication that something is happening, we need to slow down.

So don't misunderstand what I'm saying. We are absolutely not saying that we do believe something that will happen. What we are saying is that, doing this kind of change, the first time we add so many new TLDs, that cannot be done with no risk. We just need to make sure that we can handle and take care of that risk and make sure that nothing is happening.

---

The third thing -- third bullet here on contingency plans is approximately what I already was talking about.

So maybe I took it a little bit backwards. Regarding the digital archery and batching -- so what we are saying is that what mechanism that is used to decide in what order to introduce the TLDs have nothing to do with the scaling or anything. So we don't see that has any impact whatsoever on the security and stability. It might be the case that there will be an automatic ordering. We don't care. But the only thing we're saying is that there need to be some kind of feedback mechanism so this rate of which changes are made can be adjusted.

So let me stop there and start by asking if there is any SSAC member that would like to add something to what I just said.

Obviously, I don't know whether I did such a good job. But, anyways, I think it's actually a good thing to hand over immediately to a more open discussion. So, Heather, over to you.

CHAIR DRYDEN:

Thank you very much, Patrik. Are there any questions for SSAC? I see Egypt, please.

EGYPT:

Thank you. And thank you, Patrik, for the clarification. I'm trying to seek further clarification whether are there certain known parameters that should be monitored? Known to the technical people, I mean. I don't need to know them. But are there something measurable and specific that we should keep an eye on? And is it known whether the

---

feedback should mean that introduction should slow down or whether there was something that would ultimately not be stable? I mean --

PATRIK FALTSTROM:

Let me take the second part first.

As far as I remember and what we're talking about in SSAC, there is no one that have seen or can foresee any kind of feedback from this kind of system which give, as a result, that the TLDs should not be introduced. We only talk about the speed and slowing down, pause, et cetera.

We do have, for example, the dot com domain that have more than 100 million records in them, which means that having 1,000- 1 1/2 thousand records as delegations or something that we will have in the top-level domain is not a problem whatsoever.

So the only thing we need to ensure is that the whole process from IANA, when the request comes in to make a change, the whole validation of the request itself, the verification of the data of the completeness of the request data, that the request is actually correct by itself, the implementation of that in the root zone, the signing on the root zone, the publication on the root servers, we're talking about that whole system.

And, no, to go back to your first question, no, there's no single parameter you should look at here. It's more the case that we need to make sure that this system that today works smoothly -- and you can talk all of you with your ccTLD and talk to them about how sure it operates today. We need to make sure it still operates smoothly both before, during, and after the introduction of the new TLDs.

CHAIR DRYDEN: Thank you for that clarification. I have Sweden and then the EU Commission, please.

SWEDEN: Thank you very much, Madam Chair. And my question was in line with his question, so I can put it a little bit differently.

What you said before, Patrik, was also that it's a combination. When you say slow down or pace, whatever you want to call it, when there are going to be changes in the root server system, it's actually a combination of the introduction of the new gTLDs together with exchanges in the existing TLDs. So the whole -- the whole exchanging of all this old and new gTLDs has to be in a process that all of it has to be able to slow down or pause or whatever. So that's one question, if I understood it correctly.

The other question I have is about this batching system. Batching or not batching. And now we have the situation that the digital archery has been suspended. And we don't really know what kind of model or mechanism we will have for what's coming to be instead or pick up the system. We don't know.

But that also might be the case. We heard sometimes that we talk about no batching at all. And I also want -- this is also double-check, by the way, that it might be the case that one batch is just not that bad. Because you're talking more about pace and keeping -- monitoring the different parameters and having the correct pace to make sure that nothing has happened. And the third thing is actually, okay. Then I

---

have the tool maybe to slow down. And that is something I ask all the time. And I'm not sure I still have got the good answer from ICANN board that they have a tool for that. But do you think for the SSAC root that there is such a tool and, maybe we don't have to be worried about that one? Or maybe we do. Thank you.

PATRIK FALTSTROM:

So first we're going in the communitary effect. I think you're correct that there are slightly different operations that are done in this root server system when you're introducing a TLD compared to when an existing TLD wants to have changes. What you have to remember is that, as soon as you're introduced, one, a new TLD, that ends up being one of the existing TLDs because it might be the case that directly after a TLD has added a root zone, that TLD might want to change.

So, yes, you're right. That maybe is the case that now a speculating person -- this is not an SSAC view -- that, for example, there could be a priority that existing TLDs get higher priority in this -- in the workflow, whatever. It's just important that the workflow can handle both of these things that will happen.

Regarding batching and digital archery, whether one should do it in one batch or not, I think from a technical and stability standpoint, you're right that we're talking about being able to slow down. That implies that you will do things in some kind of order.

So I -- so, for example, me, I personally would rather talk about some kind of ordering. And maybe you should have some process in the cases that you need to order things. You just have to take care of that, which



is not something uncommon. For example, when you have responses to an RFP, of course they are evaluated in order. You cannot have -- as with the new TLDs, we cannot have 1,950 something people sitting parallel at the table and we agree when we turn the page and read the next page. So some kind of ordering will be there. What I'm saying is that maybe some of the ordering will happen automatically just because of the process. Maybe it doesn't have to be an explicit ordering.

Regarding a tool for slowing down, that is something we from SSAC don't have any insight in. Maybe individuals in SSAC do have insight because that has to do with the process itself of the TLD process and we're not involved in that.

CHAIR DRYDEN:

Thank you for that.

So next I have EU Commission and then Norway.

EUROPEAN COMMISSION:

Thank you, madam. Thank you, Madam Chair. And thanks to the chair of the SSAC for the presentation.

I'm afraid my question might be a little bit along with the question asked. But just for clarity, Patrik, you mentioned the need to have some kind of feedback mechanism with introduction of new gTLDs. I would like to have confirmation if you, as the chair of the SSAC can give it, does such a feedback mechanism already exist or is it going to be created? Is this a question we should ask the ICANN board maybe? Thank you.

PATRIK FALTSTROM:

I think -- what feedback mechanism is to be implemented, I think to a large degree should be -- let me answer the second part of your question first.

IANA already today do have some feedback in the form of graphs on their web page on how many items currently have in their queue, the length of the queue, et cetera. So there are some statistics there already.

Maybe the data behind the statistics is enough for a feedback loop so that ICANN staff can slow down. But I think it's a little bit dangerous for people from the outside to request a specific feedback loop. Because what data you're gathering in the form of the feedback might be different depending on, for example, how the slowing down process is. So that is something that must be part of the system itself. And that's why it's a little bit hard to talk about specific variables.

But this is also one of the reasons why, from a 10,000-meter level, I'm talking as an existing TLD owner or someone that is using the existing DNS today that a feedback loop for me, speaking personally, could be, for example, just to measure and present every day how long time does it take from a request comes to IANA to where it's available in the root zone and just to ensure that -- keep track of that value and see that it doesn't increase.

Another thing could be to look at the already existing graph that has the length of the queue of the IANA workload and make sure that the length of the queue doesn't increase. There are many things that can be done. And I think we -- at the same time as a feedback loop is needed, I really, really want everyone to concentrate and be able to

---

actually introduce the TLDs. So we should not spend -- we should spend enough time on the feedback loop but not too much. Because it's important that we still spend energy on the main loop. Was that approximately what you wanted to hear?

EUROPEAN COMMISSION: If I may, Madam Chair, yeah. It is additional information for which I thank you. But I'm still not clear whether -- why you're suggesting it. Maybe it's my misunderstanding. But I understood that you were suggesting that the feedback mechanism is necessary. But, on the other hand, I get the impression that agreed across the whole community, feedback mechanism, and what would be the proper metrics and who should perform the measurements does not exist as of today. And I'm not -- to be clear, we're not suggesting -- at this point in time, we're not suggesting that this feedback mechanism would be a necessary condition before we can make any step forward. It's simply really a direct question does it exist or does it not exist? Thank you.

PATRIK FALTSTROM: Let me first explain why we at SSAC have explicitly been talking about this this week. And I'm pretty sure that we will publish a paper with our view as a result of the discussions we have had before this week and the discussion that we have during this week. For example, next week, I think, we'll be able to say something.

Now, the reason why I stood up at the microphone yesterday regarding batching and why we at SSAC want to still talk about this and talk about the need of a feedback loop is we see too many people in the ICANN

---

community still talk about thousands a thousand a year being a hard number, and we believe that is the wrong way of looking at the problem and that's why we wanted to point this out.

So from that perspective, yes, we do see that not enough work is done on this topic. And let me now hand over to a little bit of other SSAC people who can explain a little bit, and I start with Jim Galvin.

JIM GALVIN:

Thank you. Jim Galvin, vice chair.

Let me offer the following perspective on a feedback loop. I think you can think of a feedback loop as having three parts. There are the actual metrics and the values that those metrics have. There is the second part, then, is what value you want those metrics to have, so at what point is it important to notice that there is an issue that needs to be addressed. And then the third part is what are the actions that one would take in response to those.

With respect to the first part, as Patrik said, IANA does have statistics on their Web site. There is actually a statement of a graph which shows the number of requests that are in the queue, what has been completed and what is still open. So the data is there as far as how the system is performing now. The second part would be at what point does that data require attention. And that's a discussion that should be had. It should be a known quantity, you know, what values we're trying to achieve, what trends in the values are important. And answering that question is essential.

---

And then the third part, of course, would be deciding what actions to take. And the obvious action is you want to slow down at some point. You need to say that I need to introduce new TLDs a little bit slower because I need to keep the system functioning at appropriate service levels.

Thank you.

CHAIR DRYDEN:

Thank you.

Another reply? Lyman?

PATRIK FALTSTROM:

Lyman?

LYMAN CHAPIN:

Yes, thank you. Lyman Chapin, I am a member of SSAC.

One of the things that we are grappling with right now, and one reason why we are still actively pursuing this as an open issue, is that it isn't clear exactly what part of the system is likely to suffer the greatest stress if problems arise with the rate of change.

We think of the rate of change as the rate at which we, for example, add new TLDs to the root in the New gTLD Program. But we are also adding to the root new TLDs that may behave very differently from the ones that we have in there currently.

---

The TLDs, both ccTLDs and gTLDs, that have been in the root for a long time and that we know and love have characteristics that we -- that the system has, over time, adapted to very well. And by that I mean the rate at which, for instance, a registry operator requests changes in the records that define a zone in the root.

It's very possible that some of the new TLDs that we're adding with the New gTLD Program will behave very differently. We don't know that that will be the case and we don't know how different it might be. But if you look at the list of applied-for strings and the applicants that are submitting them, you will notice that in many cases these are organizations that, in other circumstances, would be registering zones under -- at the second level or third level under existing ccTLDs or, say, dot com or dot net or something like that.

And the way in which people view reasonable change rates in, for example, dot com is very different from what we're accustomed to in the root.

So we need to be able to take into account things like that as well. It isn't just how quickly you add new TLDs to the root. It's also what kind of ongoing daily or other periodic change rate will you see from those new TLDs, which may have very different models of operation from the ones that we're accustomed to and the ones that the system has evolved very nicely to support.

So it's just to -- not to put another red flag in the air, but simply to point out that as SSAC, we are still taking this issue very seriously, and we believe that there is more work that needs to be done to be sure that we do know what all the metrics are, that we know how we're going to

---

measure them, and, as Jim has said, that we know what it is we would do if those measurements tell us that something is beginning to get out of bounds.

CHAIR DRYDEN: Thank you, Lyman. I have Norway, Germany, Egypt, and the Netherlands.

NORWAY: Yes, thank you, Madam Chair. And also, thank you to the SSAC for the presentation.

I'm not going to repeat the questions because the questions I had also was referring to some of the same, especially the ones that the commission was asking about if this feedback mechanism exists. And of course now we have had some sort of questions about this.

So I assume, then, based on the answers to this that we from the GAC side, at least, would then ask ICANN, ICANN Board, about defining these feedback mechanisms, what metrics are you going to define, and then, of course, the metrics level, and then, of course, actions based on when you'd reach these levels defined.

That is something that I think we, as the GAC, can ask ICANN Board about.

Also, one question regarding being, also, I think what Egypt was asking regarding parameters. Are there any -- In addition to what you have said, and this was sort of in the processes for updating the root zone file, implying, well, what VeriSign are doing and what IANA is doing and

---

so on. Are there any parameters that the root server operators should monitor? And then based on any sort of things there, are they -- Does that have anything to do with this or not?

That is one question.

And, also, just a last comment. My question was also regarding what Patrik was saying, that I have also heard, that this thousand a year would be a thing. So the conclusion here, at least what you have said, is that the number one batch or not does not really have any -- it's not a problematic to have more than a thousand, then, as long as you have a sort of controlled rate of introducing it into the root.

Thanks.

PATRIK FALTSTROM:

Yeah, regarding the root server operators, I think the various reports that are written and if I remember correctly what RSSAC said, it said that all the root server operators have no problem at all to handle this and they don't see any problems with it. But that questions needs to be directed to RSSAC, not SSAC, I think, is more proper.

CHAIR DRYDEN:

Thank you for that, Patrik.

So next I have Germany.

GERMANY:

Yes, thank you, Madam Chair. I think some of my questions have been answered already during the discussions, but I just want to summarize



---

what I have gained from this discussion. And maybe you can correct me if I am wrong with my kind of understanding.

I now understand, because I see relationship between batching and root zone scaling from the other way around, I think. As far as there are problems -- or unsolved problems in the root zone scaling, the more important it is to have a neutral batching system or a neutral system for setting priorities.

If there were no such restrictions, I think it would be easy to resolve the problem. But as far as there are uncertainties, we may face the question of having fair chances for every applicant to be at the beginning of the queue.

And I learned that there are still some uncertainties in respect of the future steps. You can probably begin with some -- with the numbers -- the phone number of applications, but then you need some reaction, and you need to see whether there is a need for some kind of interference.

And this brings me back to a general -- maybe it's too easy but the general time planning for the entire project, because we are now in a situation that probably at the middle of next year, there will be the first applications ready for delegation. Then it sets a very optimistic point of view. Then it would last, let's say, around about two years, in this optimistic point of view, when the last delegations would be possible. And that means that you would have, let's say, two years for these 2,000 delegations. As long as there are no necessities to react because you see that there are unexpected interferences from -- after the delegation.

---

And, yes, is this kind of understanding correct?

CHAIR DRYDEN: Would you like to reply?

PATRIK FALTSTROM: I think that it's something I'm happy to discuss this off-line with you, but I think one of the important things to remember is as soon as we start to add things, the more we get to know how the changes are absorbed. And, for example, what Lyman explained, that today we have between one and two changes per year per existing top-level domain, and we don't really know, for example, whether that number will change.

So there are so many uncertainties now that I think many people, including me personally, think will be much, much, much clearer as soon as this system actually starts.

So today, we are guessing a lot, and that's why we -- we, from a technical community, wanted to stop guessing and, instead, talk about this system that was needed to be able to ensure that the whole process is still stable. We thought that was much more convenient and talk about that.

Anyone from SSAC that would like to expound on this?

CHAIR DRYDEN: Thank you, and thank you for the question, Germany.

So next I have Egypt, then the Netherlands.

EGYPT:

Thank you. And again, it's not a question, per se, but, rather, sharing my understanding, as Germany.

What I understood is that now "batching" is not the right word to use. We are talking about two phases, basically, the evaluation and the delegation, so maybe batching could be a used term for the evaluation phase, but then when we speak about delegation, I think if we are going to slow down, as you said, then we are going to slow down whether the domains belongs to the same batch or not.

So it's -- I mean, there is no need to know the end of batch and start of next batch, but, rather, the ordering of the delegation, which I understand would be affected by two factors: the administrative cycle and the technical performance. The administrative cycle, which could be estimated somehow, and the technical performance that has to be monitored real time; right?

PATRIK FALTSTROM:

Let me say that what SSAC are talking about is the technical work that must be done when the application sort of, quote, unquote, "reach IANA" and it start to get implemented into the root server system. That part, batching doesn't matter. It's the ordering and the ability to slow down.

There are previous parts in the whole new gTLD process that has to do with the evaluation, the various evaluation teams and various other kind of things that we at SSAC absolutely don't know even close to as much as you do. That part might need batching.

---

So I think it's a little bit dangerous to -- I'm absolutely not saying that and it's not SSAC view that the word "batching" should not be used in the gTLD process. I am just saying that for our part, what we are nervous about, we are looking instead at ordering and the speed and not batching.

CHAIR DRYDEN:

Thank you.

Okay. We are getting close to time so I have a request from the Netherlands. Is that Italy? Okay. And then I'll move to close. And I see Pakistan, and then I will move to close this session.

So Netherlands, please.

NETHERLANDS:

Yes, thank you, Heather. And thank you, Patrik, for the exposé. A couple of questions have already answered, and also due to other colleagues' questions.

One thing which I'm still a little puzzled is about who does what. And I think we have also, Heather, the report on the root zone scaling from ICANN, which says that, okay, there are a lot of monitoring mechanisms in place and different organizations working on monitoring.

Only my real question is in this feedback loop which you mentioned, where for which still the metrics and let's say the pyramids have to be designed and, let's say, choosing the right pyramid -- I think you mentioned, also, the lead time which it takes to have a change being done, et cetera, my worry is a little bit who does what? Because there

---

are so many parts involved. You refer to the RSSAC. We would be -- We would like to know if in this feedback loop, which will give information back to slow down the rate of implementation, what is this chain like? Which part you will do what?

And maybe it's not addressed to you, exactly. Maybe it's addressed to ICANN. But still, I'm a little bit puzzled about where this is designed for this early warning monitoring system and the parameters, and which parties are involved.

Thank you.

JIM GALVIN:

So, Jim Galvin. I would respond by saying that we have asked the same question. I mean, asking for an early warning as part of the response to the root scaling report, you know, that is, in fact, a question that does need to be answered.

We're focused here on saying that a feedback loop should exist, and that's what has to be created and understood and documented so that we can all be comfortable that we can see what's happening with the system and that it will be reacted to appropriately.

So we don't have a specific proposal on what that should be, but that is something that probably we should look to staff to develop and make visible for comment by the community.

PATRIK FALTSTROM:

As you pointed out, there is this report that was just -- that I do know that you got the other day that explained a lot of these things are

---

measured, and this is also why, as a response to the European Commission's question, I do not think that we -- it might be the case that we don't have to invent any new measurements or new variables. We just need to make sure, just like you're asking for, we just need to lay the puzzle and see what we're doing when various things are happening.

CHAIR DRYDEN:

Thank you.

So next, we have Pakistan, please.

PAKISTAN:

Thanks to SSAC for useful updates on matters under discussion. As we all know, that stack matters are important and ongoing activities which are mandated for the successful Domain Name System. So I request SSAC may update the GAC the solid and fruitful input from academia on it, if still there's a lag from academia side. Moreover, I just suggested that SSAC may get feasible recommendations from the academia through providing their input by performing R & D work by maybe they offer a scholarship (indiscernible) scholarship to the concerned and technical teams.

Moreover, keeping in view the discussion matters especially on batching and digital archery, it is suggested that ICANN may continue multistakeholder approach in order to find out the feasible solution to these issues about the processing of applications for new gTLDs with the support of ICANN supporting organization and advisory committee, including GAC. In this regard ICANN also consult with the global

---

community as well. Definitely ICANN require adequate time and full support of all relevant stakeholder. I hope through this -- eruption of this approach, ICANN will find out the possible solutions to these issues. Thank you.

CHAIR DRYDEN:

Thank you, Pakistan.

I skipped over Italy earlier. With apologies, please, Italy.

ITALY:

Thank you, Chair.

So what we understood from this discussion is that, first of all, 1,000 new gTLDs per year is a number that should not give problems, but provided that, of course, that they are not added all in the same day. That has to be made progressively in the....

But what is real important from this discussion is that based on a continuous monitoring of the performance of the root zone and just in case there are some degradation of the service, then there is a contingency plan and we have to be assured that this is there. And then it is important who is doing what. So that is -- and there would be -- and there should be an impact on the activity of the ICANN staff and, consequently, the Board in order eventually to make some retard in the insertion of new gTLDs. But what is really important is to perceive that the process has a continuous monitoring and that there are good enough contingency plans.

---

So this was -- I would like to convey my compliments also to the SSAC because we get the impression that there are a lot of variables, but in the same times, also a good knowledge of all the variables and their possible affects.

Thank you.

JIM GALVIN:

Thank you. So Jim Galvin. I have been asked to clarify a comment that I had made earlier which relates to what Italy was saying and others have said. I had perhaps overstated -- or I did overstate when I said staff should be asked to create, you know, the early warning system and monitoring system and it's important to emphasize that much of this activity is taking place already and it's not really a staff activity. You know, the Root Server operators obviously already extensively monitor their systems and so they're very well aware of the function of the root system and how it's working. There's also the SSR activity which has been active, using the community to develop the metrics to look at the DNS and understand its operation and in fact there will be discussions later this week on Thursday where they'll be talking about their work and the results there. The role of staff in all of this is simply to support and facilitate the community and the Root Server operators and also to facilitate the community continuing the development of those metrics, what those metrics mean, and the contingency plans of what to do according to the values of those metrics.

So I just wanted to clarify, it's not really a quest of staff to do the work but for the community to do that and also to add to the work which is already in progress. Thank you.



CHAIR DRYDEN:

Thank you, Jim. So with that, I think we can conclude this particular session. And with thanks again to the SSAC for coming for what is a very interesting and I think timely discussion about issues related to the root zone and delegations and monitoring of those delegations.

Okay. So thank you. And hopefully we'll meet again at the next meeting as well.