DUBLIN – Fellowship Morning Sessions
Wednesday, October 21, 2015 – 07:30 to 09:00 IST
ICANN54 | Dublin, Ireland

JANICE DOUMA LANGE: When you come in, come on down. Patrik does not bite. Been around a lot.

Good morning, interpreters. I love how I say, "Good morning, interpreters," and everybody here says, "Good morning."

Headsets for anyone. Remember it's not just about you speaking another language. It's about you understanding someone else's. Headsets are in the back. Unless everybody understands Russian and Arabic, you might want to get them. I know I'm failing miserably in those two languages as hard as I try. Spanish as well. Headsets are that thing you want to pick up on your way in the door.

I'll just patter on a little bit. I know you'll be right back. Wednesday… Today we have, at the end of the day, another one of our accountability sessions. If you have been studying, if you've been not able to attend, this afternoon is your best time to spend a couple hours with the working groups and the experts who have been working on this for months and hoping to come to a grand conclusion for all of this at this meeting. Today is a good day for that.

We have, I know, an IDN session today. Nabil, you are presenting today, if Nabil [inaudible] 12:00? 1:00 P.M. So in support of Nabil's presentation as well, if you haven't had yourself signed up to another session, I know I'm going to try to be there at 1:00 in support of our alumni in other session.

This morning, we are going to start off the morning with Patrik Faltstrom, who is the Chair for the Security and Stability Advisory Committee.

We have the possibility, up there to my people who would love some free time today, of not being here after this. We had one presenter last night who had to cancel, and another one is not sure because of the workload. So that's kind of a win. It's all right. We'll see how that goes, but for right now, Patrik, good morning. It's yours.

PATRIK FALTSTROM: Good morning. Good morning, everyone. I heard that you are talking a little bit about the IANA transition issues. I have some good news. As you know, the transition itself is done by a group called ICG, which I also happen to be a Co-Chair of, requested from the names, numbers, and the protocol parameter operation communities to come up with the proposals that we would compile into the solution for how to [continue] to run IANA.

We have not met with you in ICG, and that's just because our work is going quite smoothly. We actually had an open session Monday here in the auditorium where anyone could ask questions if they wanted. We got zero questions, believe it or not. So we closed after five minutes, so people got 55 minutes back.

But then the people dealing with accountability, which I have no idea, honestly, why they are talking so much about accountability because the IANA transition is not much accountability at all. People are talking about what they're unhappy with ICANN, which is a completely different thing. To some degree, I'm a little bit irritated over that group not really keeping to their [stay and] to their charter. But on the other hand, the chairs of that, they're doing the best they can, and I wish them all the best.

I'm the Chair of SSAC, and that's why I'm here. Can we get the picture up again? First slide, from the beginning. Thank you. Forward. And forward. There.

So who are SSAC? SSAC is the Security and Stability Advisory Committee over here in ICANN. Our charter is to advise the

ICANN Board and the community on anything that is security related or anything that has to do with the identifiers that we're using.

We have created 73 recommendations. There are 73 documents with various different kinds of recommendations so far.

Next slide, please. We are 35 members. Our charter that we are to advise the ICANN Board and the community matches the mission and core values of ICANN, which you can see up to the far left if you have very good eyes: to ensure the stable and secure operation of the Internet's unique identifier systems, and preserving and enhancing the operational stability, reliability, security, and global interoperability of the Internet. That is what ICANN must do. This is part of ICANN's goal.

Our charter is to advise the ICANN Board and community, so we are the ones that try to ensure that ICANN's community and board is actually following; they're living up to their mission.

So that's what we're trying to do. Not always easy.

How do we do this? Well, we come up with these reports, and these reports include – some of them – recommendations. Some of these recommendations are recommendations to the ICANN Board, and you'll see that to the right, that we send advice to the board and the board then decides either that they have to include that in some kind of policy – they have to inject what we're recommending in the policy development processes. So we might say, "You should now think about these and these and these things," and they have to start a PDP. Or it might be the case it is something that ICANN staff should do better or different, and then the board asks staff to, "Please ensure that you are doing whatever SSAC is proposing."

Or it might be the case that we ask ICANN to please talk to the W3C, or the CA Browser Forum, or the IETF, or some other organization, ITU, and the ICANN is doing that. Or the fourth alternative, which is very important. ICANN Board makes a decision that, "Oh, ha ha, you SSAC people. Yes, you talk about risk, but are you prepared with taking that risk? So we are choosing a different path forward than you suggest." So to take advice into account does not mean that they have to follow our advice all the time, although they have followed our advice, I think, in all cases except two.

One is a little bit older, but one where they did not follow our advice had to do with namespace collision, where they did choose a different path forward than what we choose. We suggested that ICANN should set up as server that actually received traffic that was coming to all new TLDs so that it would be possible to analyze where the traffic was coming from. ICANN decided to not do that, but instead force the registries to delegate all new domain names to 127.0.53.53 and not collect data.

The final report from ICANN that explains why they did choose a different track is not published yet. It's not even we know why ICANN did choose a different path. But of course, we in SSAC chose our normal escalation methods when the board does not follow our advice, which was we printed T-shirts which said 127.0.53.53 on the back. That's how we escalate.

Anyways, we have good cooperation with the board, and we had a meeting with the ICANN Board yesterday evening. So everything is fine and dandy.

But it's important to know that there might be other parameters that makes it reasonable for ICANN or anybody else to not follow our advice, but they must take it into account, which is the wording.

Next slide. I think we are – yeah. Good. This is last thing. What could be interesting is the box down to the right, which is outreach. It's our webpage. We have a Facebook page, new for us. Don't really know how to use that. Modern things.

Also, we have started to make videos to explain what we are doing, to explain our reports. We are very modern, start using video now.

Okay. Anyways, is there anyone who has any questions about SSAC? Because I have, thanks to [inaudible], some questions from you that I was thinking of going through here.

Nothing special? Okay. Shoot. 17 minutes. Then we have the Chair of RSSAC, which is already here. Why are you here to early?

LARS-JOHAN LIMAN: Because I didn't upload my presentation in time.

PATRIK FALTSTROM: Okay. Okay. He didn't upload. Okay, you can go to Yannis and do that now because I don't need it anymore. So I'm fine. Thank you.

Lars-Johan Liman, Chair of RSSAC, which is sort of a sibling organization to us. He will talk to you shortly.

First question. How does IDN effect the security of the DNS?

Well, being the person that came up with IDN – I'm the person that wrote the standard – it's my fault if you think it doesn't work really well. Blame me. Yeah. Okay. Sorry.

Anyways, no, it was me and a couple of other people. So more seriously, IDN do not affect the DNS at all because, as those of you who work with IDN know, the IDN or the use of Unicode or something that has only happened in the user presentation, the clients in the presentation layer. Because every time an IDN is really [used in] DNS, it is first converted to what is called an A-label or only ASCII characters, A-Z, 0-9, and dash. Because of that, DNS only sees those ASCII characters.

This was actually one of the main reasons why I proposed to have this encoding and not start to parse around Unicode characters in the DNS, because it was too risky to start to use Unicode characters in the DNS software we have out there in the world. DNS does not use Unicode. Okay?

On the other hand, all provisioning software, all web browsers, all e-mail clients, all e-mail servers, everything that supports IDN must support the transformation between what is called the U-label and A-label. That's where there are problems. That's where there might be bugs. There might be clients that cannot do the transformation, so when you get an e-mail or you see a webpage, you still see "XM," dash-dash, blah-blah-blah, because it's not transformed from the ASCII that is used in DNS to Unicode.

This choice was made just because, one, it was a little bit risky to start using Unicode in the DNS, even though the DNS protocol theoretically should be able to handle it, but also because there are so many clients that could not handle Unicode. The most important thing we felt was to be able to do a reply on an e-mail, and that is only possible if it is the case that it just received the ASCII version. So inside the system, only ASCII is in use.

UNIDENTIFIED FEMALE: Hello, sorry. I really want to understand this, but it's a bit hard for me when you use the acronyms. So if you can repeat this is a more simplified way, I'd love that. Please.

PATRIK FALTSTROM: Okay. [inaudible]

UNIDENTIFIED FEMALE: That's okay. [inaudible] Not a question, but also, can you make examples so that it's not technical sound, those who don't come with a technical background could understand something?

PATRIK FALTSTROM: Just a second. Sure. Okay. So if you want to go to a webpage which has a domain name with a non-ASCII character in the domain, you type in that string in the web browser. The web browser detects that, at least one of the characters in one of the tokens of the domain name, where one token is a word between dots. So for at least one of them you have at least one character that is not ASCII.

What the browser is – yeah. What is ASCII [inaudible] yes, which is not A-Z, 0-9, dash, and a few other characters. What the browser is then doing is converting that into XM dash-dash and couple of other characters.

Okay, wait a second. I need my own computer. My own computer, which is not here. It's down there. Anyway, this is a little bit early in the morning. Now I want to actually show explicitly how it works, and I need software that I have on my computer. So let's ignore that now.

What is happening is that the browser is converting that specific token into something which is without Unicode and only in A-Z and 0-9, the letters we use in Western Europe and some other languages in the world.

That string is what is used for the DNS lookup. So the string that includes the Unicode character is not what is used in lookup, and that is the string that is used in the URL when the browser is actually fetching whatever resource it is to fetch.

UNIDENTIFIED MALE: Version 2.

PATRIK FALTSTROM: Yeah. Yeah. So there are a multitude of websites where you can type in Unicode strings and get back this string which does not have Unicode character outside A-Z, 0-9. When a string includes Unicode characters and is a domain name, it's called U-label. When that string is converted to ASCII only – A-Z, 0-9 – it's called A-label.

The new version of the internationalized domain name, which we are using nowadays since 2010, has a one-to-one mapping between A-labels and U-labels, and earlier versions did not have that ability to map back and forth.

Nowadays, if you use more modern software, whether you use a string with the Unicode character, or whether you use the XM dash-dash version doesn't matter. A-label and U-label. Nothing else.

Unfortunately, the program here at ICANN for universal acceptance, they had in their presentation Monday many of the kind of wordings on these kinds of things, which is actually a bit confusing. A-label, U-label – nothing else. It's really important to use the correct terminology.

Okay. Ten minutes left. I must move forward. We are to do IDN some other time. Yes? You wanted to add something?

UNIDENTIFIED MALE: Just about IDNs. [inaudible] from Morocco. I'm here as a returning fellow, and I'm also a member of the Task Force for Arabic Script IDNs. So we have come up with the first proposal from the community and still have just some comments to update our final proposal to ICANN.

My question is, our items in the timeline is to tackle the second level for the top-level domain. So just the second level, so we are now working on the top-level domain, the Arabic script to be used. The second step will the second level, and the third step will be the universal acceptance, which means how to deal with e-mail addresses and how to have e-mail addresses with IDNs. What are your suggestions to complete these tasks?

PATRIK FALTSTROM: Well, I think, first of all, ICANN should create a good – well, as we had said in SSAC, it's important that ICANN come up with one set of code points and rules for those code points to be used in the root zone. ICANN don't have to do anything else. That is the important thing. What is done in the TLDs are up to the TLD registries. Of course, ICANN can come up with suggestions to help coordinate between the registries, but that's absolutely secondary.

Then ICANN should work on universal acceptance. So from my perspective, concentrate on the root zone, and then look at the universal acceptance. Second level is not important, not from an ICANN perspective, because ICANN don't have any policy [ready] for the second level.

So coordination, yes. But that's also why there's a big difference between the root zone and the second level.

UNIDENTIFIED MALE: [inaudible]

PATRIK FALTSTROM: E-mail? That's part of the universal acceptance. But that is also not hard because there are three or four dominant producers of e-mail clients. Just talk to them.

Regarding the Arabic script, personally it's cause of confusion between me and ICANN staff. My comments on the Arabic and Armenian LGR – [someone] received that only the day before yesterday. I'm sorry. But there are some issues there.

UNIDENITIFED MALE: [inaudible]

PATRIK FALTSTROM: Someone got it two days ago. Excuse me?

UNIDENTIFIED MALE: Today, too.

UNIDENTIFIED FEMALE: [inaudible] microphone.

PATRIK FALSTROM: I don’t understand. He said, “Today, too.” I don’t understand that. Anyways.

UNIDENTIFIED MALE: [inaudible]

PATRIK FALTSTROM: Yeah. But I’m not talking about the workshops. I’m saying that I sent in my comments on the Arabic and Armenian LGR proposal two days ago only. My apologies. That’s what I’m trying to say.

So next question. We have six questions –

UNIDENTIFIED MALE: I have a question.

PATRIK FALTSTROM: No, no, no, no. I have my questions, and those the ones that I’m going through first. If you still have time and Liman has not [thrown away]… Okay. I’ll do this fast.

"What are your ideas about how to encourage ISPs to enable DNSSEC at their site? Do you they have to wait for a disaster and then deploy DNSSEC, etc.?" And the answer is, no they have not and they should not wait for disaster.

ISPs, or more importantly the parties that look up domain names on behalf of end users, which often are ISPs, they should and must turn on validation or DNSSEC-signed responses. That's what they should do.

Before validation existed, people would not sign their zones. So please talk to your ISP, your enterprise, or whatever – your IT department – to turn on validation. It doesn't cost anything. It doesn't cost CPU. There are a lot of rumors out there. Just turn it on.

Next question: "DNSSEC and then DANE? What is next?" That's actually a damn good question. IPv6 deployment, DNSSEC deployment, and DANE deployment is moving slower than a snail. I talked to Dan York yesterday at ISOC that is running their Deploy360 program and asked him what we will do? I asked this question to him: "When you are done with IPv6, DANE, and DNSSEC, what will you do?" He just looked at me, like, "That would never happen, that we're done with all those things."

So we agreed that when IPv6, DNSSEC, and DANE is deployed, we, and also all in this room, can retire. More seriously, we don't really know. But it's so much work to do with DNSSEC and DANE, and DANE is something we need because we cannot use certificate authorities anymore for certificates. It just does not work. Okay. We must use DANE.

For those who don't know, DANE implies that you create your own certificate, place it in your own DNS zone, and you sign your zone with DNSSEC. If you do that and software supports it, you do not have to go to a certificate authority and pay for certificates. You can manage the certificates yourself. That's the whole idea.

Of course, the certificate authorities don't like that because that means that they lose business. But I don't really care. I want to [inaudible] the world.

Next question: "Openness means more security needed. Could you share with us your ideas about that?" Yes. Openness means that we share more information between us, between services. You log into one service and then you use that log-in when you go to a second service, etc. That implies that these services have a higher responsibility to keep track of your user name, your password. The log-ins need to be much more secure, and people are not allowed to – it's a stronger responsibility to take care of whatever kind of mechanism that I used for specific logging in and information about persons.

We in SSAC call that credential management, or credential lifecycle management. We are currently on a report on credential management lifecycle that hopefully will be released in just two weeks. Whoever was interested and asked this question, there will be a report in the next two weeks, and we'll present it at the Internet Governance Forum in Brazil. The idea was to have it ready to this meeting, but we just didn't make it. That was a second apology. Sorry.

By the way, in the slide deck that Yannis has, there is a presentation on this credential management in that slide deck, some slides that you can have a look at.

Next questions – three minutes. Three minutes, yeah. "There are various aspects of securing domain names and IP addresses, both for providers and end users. What are your activities for both groups?" Well, on the domain name side, we in SSAC are working with DNSSEC and DANE on namespace collision and various kinds of things.

On the IP address side, we are following the work that is done with, for example, RPKI, but we don't see the world really agreeing on the use of our RPKI yet, which is signing of IP addresses and route announcements.

So what we're doing on the IP address side is that we have currently started investigating what IP addresses people are using on the Internet as compared to what IP addresses are allocated for people on the Internet. Now, when we have run out of IP addresses, people have started to use IP addresses that are allocated f0r someone else, but not announced. That is something that we're currently investigating because that has increased.

So if Yannis, for example, I got a block from of IP addresses from AfriNIC and she is not using it, then Lars-Johan Liman might detect that and start using those addresses.

UNIDENTIFIED MALE: Stealing them.

PATRIK FALTSTROM: Yeah, stealing them. Yeah. But on the other hand, as long as there's no overlap, as long as not both Yannis and Liman use the addresses, it is sort of okay, although we don't know who is using them, so it ends up being problematic for law enforcement and stuff and things. So that's what we are doing, and the first step is to look at this.

The next step, of course, which is worse, is if all IP addresses are not allocated but if all IP addresses are used, because I envision and even SSAC envisions that at that point in time, people will start to use IP addresses that others are using as well. That's really bad, and that is unfortunately specifically people not living close to the core of the Internet in the world, which means that specifically developing countries will probably be the ones that have the most issues in those cases. Or the northern part of Sweden for that matter, which is pretty far away.

At that point in time, we must use IPv6. So IPv4 would probably I think personally be more or less unusable within the next two years. That means that everything you are installing today that you think that you will use for more than two years must support IPv6. So even if people like aid organizations from Sweden or whatever donate things to do and it doesn't support IPv6, do not accept it because it will not help you so much.

Of course, it will help for a short period of time – the next year or something – but you must be aware. Put a special sticker. Mark the equipment or software so you know, "I have received this. I did a conscious choice to use it, even though it doesn't support IPv6." Right it down on the list so that you, one year from now, replace it. You cannot use IPv4 much longer than one or two years.

Then the last thing: "DNSSEC is a solution that should be utilized by the ISPs to ensure the originality of the domain IP conversation and avoid attacks like DNS poisoning. But how can you help end users to protect themselves against such attacks if their ISP is not protecting themselves by tools like DNSSEC?"

I already talked about this. It's actually sort of confuses a little bit, securing the IP addresses routing with the domain names. How can you protect yourself? Well, you can use your own resolver. You don't have to use the one that your ISP is using. You can use the Google 8.8.8.8 resolver that is doing validation of DNSSEC. So you don't have to use the resolver that the ISP is using. That's probably the easiest way.

Okay. Last question.

UNIDENTIFIED MALE: Is it secure to use public Google DNS for a resolver?

PATRIK FALTSTROM: Yes and no. Whenever you have a resolver that you communicate with, the communication between your software and the resolver is unsecure, so you must be sort of certain that that part is trusted, regardless of whether it's between you and the Google resolver or you and the resolver of the ISP. The best thing is if you run a resolver on your own computer, of course.

[MANUEL]: Hi. [Manuel] form Mexico for the record. I was wondering why is SSAC an AC and not an SO for doing policy?

PATRIK FALTSTROM: Exactly because of what you just said. We don't run a policy development process. We're an advisor to the policy development processes.

[MANUEL]: But given the fact that what you do is very, very important, shouldn't it be better to make it policy creation instead of just giving advice that might not be mandatory?

PATRIK FALTSTROM: Yeah. No. We actually think it's much, much better if we have a strong advisory role, which means that we are forcing the policies development by others to have certain sorts of features.

For example, we were the ones that forced all accredited registries and registrars under the new RAA 2013, and all new gTLD registries, to both support IPv6 and DNSSEC. So we were able to influence that policy, even though we didn't develop the policy. That's how we work.

UNIDENTIFIED MALE: Yeah. Sorry. After the transition from IPv4 to IPv6, this –

PATRIK FALTSTROM: No, no. It's not a transition. What we are doing is that we are deploying Ipv6 while still running IPv4. Then there's a separate step to potentially one day turn off Ipv4.

UNIDENTIFIED MALE: Okay. Ipv6 is 128-bit, so the message after DNSSEC signing our zone will be bigger.

PATRIK FALTSTROM: That is correct.

UNIDENTIFIED MALE: This means the [inaudible] of 1,500 will be blocked by a firewall at our side.

PATRIK FALTSTROM: Well, okay. What he's saying is the packet size for the DNS response will go over 1,500 bytes. Yes, that might very well be the case. If it is the case that you have a firewall that blocks IP packets that are larger than 1,500 bytes, then you should reconfigure a firewall. What you have to remember here is that

ICANN packets larger than the MTU size will be fragmented, and the MTU size is normally 1,500 bytes.

But you have to differ between the packet size of 1,500 bytes and the IP packet size, which might be larger. So what you might talk about is that your firewall is blocking fragmented IP packets. So even if you have a MTU that is 1,500 bytes or smaller, you can still handle larger IP packets by getting fragments.

Yes, there are firewalls that do not handle fragmented IP packets, specifically fragmented UDP packets. That is the problem. I agree with you there. So the only thing you can do there is in that case is to watch very carefully how large the packets are when you are signing your zone and choose, for example, the domain names for your records carefully so that compression that can be used for DNS to make each one of the records in the resource record set as small as possible so that, for example, by choosing different naming, you can get smaller packet size.

UNIDENTIFIED FEMALE: Hi. My name is [inaudible]. I'm from [inaudible]. This is my first ICANN meeting. My question is, what is the security issues that can be taken into account when deploying IPv6 within DNS? I mean due to the differences between the structures of IPv6 and IPv4.

PATRIK FALTSTROM: I think when you start deploying IPv6 networks, you should take advantage of the large address size. For example, with IPv6, you are not forced to do subnetting and the [inaudible] network according to some kind of limitation because you have a very small IPv4 address space. So you can actually, with IPv6, build a logical networking structure that matches your needs. By doing that, you can build a more secure environment where you can control the communication between the various functions in your organization. That is one of the more important things you can do.

Another thing you can do is, of course, that I find it being, for security reasons, a good thing that everything on the network do have global unique addresses. Some people think that network address translation boxes are a good thing, but actually I think that from an administrative perspective it is not good with NAT boxes because when you just see an IP address, it's so difficult to keep track of what it is. With IPv6, everything can have its own address, which also increases security.

Yeah, we need to give the microphone to [inaudible]

RAPID SUN: Good morning. My name is Rapid from Cambodia. Usually when I go to some country when they block the access to the Internet, I change to the Google DNS. As you said, the connection from us to the Google server is not secure. Can SSAC set up its own [inaudible] DNS server to make it secure?

PATRIK FALTSTROM: Well, if it is the case that you're behind some kind of firewall that are blocking traffic, regardless of whether it is a state that blocks it or a company that blocks it – hotels sometimes do all different kinds of blocking – sometimes it's easier to break out of that by having more services on your local computers. Running your own resolver sometimes solves it, so you run your own DNS server.

Sometimes you unfortunately must have a server somewhere in the world, and you open a tunnel to that server, and you tunnel everything there. There are servers to buy. For example, I have one that I bought in the U.S., a virtual server that costs me ten U.S. dollars a month. Just by having that one and open an SSH tunnel or a tunnel of HTTP just to that server, and then I tunnel all IP packets to that server. On top of that tunnel, I can then do whatever I want.

So two ways of getting around these kinds of policy things, but we have to remember there is a reason why there is a policy there, which means that what you are starting to do, what we are talking about now, is violating the policy that someone has applied for some reason. That is something I should be conscious of.

But given that you would like to try to violate that policy for some reason – and there are multiple reasons for that; I do it all the time, including my own office – you either have more services on your laptop instead of using the ISP, or you have more services on another machine that you control and a tunnel. Normally, you need to have both of these sort of – I'll call it toolboxes, just like a carpenter that carries around the toolboxes. Unfortunately, that is needed.

Thank you very much.

JANICE DOUMA LANGE: I shudder to think of your office. Patrik, thank you. A couple thoughts. One, interpreters, thank you. I know in the past trying to tell Patrik to slow down is like trying to step in front of a freight train. His passion overruns, and I really appreciate your keeping up.

Two, one of our fellows, [Elsa], mentioned to me yesterday in the download session that we should have an ICANN for Dummies, a Newcomer for Dummies, because it is really difficult, not matter how we try. ICANN Learn: I talked to her about how we would get more courses that kind of get to it.

I thought about this with this because the interest is here, and the level, as I explained to everyone, has to be set around the middle. I'm challenged. Raise my hand. Or this is not challenging enough, raise my hand. And bring it from the middle. But perhaps we can work with you on something that's a DNSSEC and RSSAC for beginners.

PATRIK FALTSTROM: Yeah, DNSSEC is easy because – unfortunately you missed that session, but we in SSAC always have DNSSEC for Beginners on Monday afternoon, which you missed. A little bit more advanced DNSSEC session starts now at 9:00 and runs until 2:00 P.M. I think, a whole day just for DNSSEC. So DNSSEC for Beginners we do have, and I do know that the IDN people have similar programs, but not by us from SSAC.

JANICE DOUMA LANGE: Thank you very much. Liman, I need to jump up and get your presentation going, so I'll let you introduce yourself while I do so.

LARS-JOHAN LIMAN: Hello. Now we're all worn out the security-related stuff. I'll put this here so I have it nearby if I need it.

Thank you, Patrik. Patrik is an old, longtime friend of mine. It so actually happens that we work for the same company. I'm Lars-Johan Liman. I'm one of the two Co-Chairs of the Root Server System Advisory Committee. I was instructed to forward my second Co-Chair, Tripti Sinha's, regrets that she couldn't be here. She unfortunately has to leave this morning because she has an important family event to attend back in the United States. Actually, her niece is getting married, so she has a wedding to attend.

But I'm here. I will try to tell you some things about the Root Server System Advisory Committee, what we do, what we're here for, and what's going on now.

Next slide, please. Can I see this slide somewhere? I can here I suppose. Yes. So what's RSSAC? Well, the root server system, as you probably know, is the entry point to the DNS system. When you want to look something up and you have no idea where to start, the root servers are where you start.

There are 13 IPv4 addresses and I believe 11 IPv6 addresses that you can send your query to to get into the DNS system and get referrals so that you can work your way down through the system and eventually obtain the information you're looking for.

The Root Server System Advisory Committee is another advisory committee within ICANN. As Patrik said, we're a sibling organization, but our scope is very, very narrow. What we do is advise the ICANN community and the ICANN Board on matters relating to the operation, administration, security, an integrity of the Internet root server system. That's very narrow. We only care about the root servers and the root zone, which is the data that we provide from the root servers.

I should mention that I work for Netnod, and Netnod operates one of the root servers. They are denoted by letters, so we have a.rootservers.net, b, c, d, and in my case, I operate i.rootservers.net together with the Netnod staff.

What we do from these root servers, which are ordinary DNS servers – there is nothing strange with a root name server. We use exactly the same hardware, exactly the same operating system, exactly the same software for providing DNS service, and exactly the same type of data in the databases. Ask any TLD or any corporation or any association that runs its own DNS service. There is no magic.

The only thing is that we happen to sit on the IP addresses that are the entry point to the system.

So we care about the root zone, which is the zone file, the list of DNS records, that we provide from the root servers. We need to have that copied into the machines, and we receive it from – the root server operators receive that from the IANA, which is operated by ICANN, and then through a process where it's approved by the NTIA, which is a subdivision of the U.S. government, and then actually provisioned by the company VeriSign. That's how it's all set up. So Netnod gets the information from VeriSign.

The entire process of making sure that the data is okay all the way from when it's generated at the IANA until it reaches the root server that sits on the network and provides the answer, that path is what RSSAC cares about.

We communicate on matters relating to the operation of the root servers, and we do that a lot with the technical community because what we worry about is the stability of the system, and the resilience, and the reachability.

Our first and foremost priority is to make sure that you get answers when you ask a root name server, and you should get the correct answer as long as that's possible. As Patrik mentioned, when you send a packet across the Internet, you have no idea how it travels. But I can promise from the Netnod side that, if your packet reaches us, we respond with the correct information as we receive it from VeriSign, and when we send the packet out with the response, it's okay.

Someone may change this as the packet goes back towards you. We cannot do anything about that. But when it leaves us, it's okay. And if you ask for it, it's signed with DNSSEC. This is a flag in the query. You can say, "I want DNSSEC," or, "I don't want DNSSEC," but if you do set that flag, it's signed and you can verify the content.

We also communicate on the administration of the root zone. We worry – I shouldn't say, "Worry." We analyze and keep track of things like the data bits in the zone file. There are more things than just domain names. There are time-to-lives, there is delegation information and stuff like that, and we try to keep an eye on how all these things are tied together so that it will function in a stable and good way.

Next slide, please. We also engage in threat assessment. The root server system is a target for various types of attacks. I will not deny that. We constantly monitor the system, the root server operators monitor, and we talk about that in RSSAC and also in other [constellations].

We also respond to requests for information or advice from the board. Later today in this room, RSSAC and the board will have a joint meeting, where we will inform the board and respond to some queries. At least it's an open meeting, so you're welcome to attend.

We also make policy recommendation to the ICANN community and to the board, much in the same way as SSAC. If we see that a supporting organization is developing a policy that may have a bad effect on the root servers – maybe they're trying to construct something that will have a very sudden impact on the service – then it's difficult for the root server operators to predict how that will affect the service. So we may forward a recommendation saying, "Don't do this very quickly. Do it slowly so that we can follow the system and see that it continues to operate correctly."

We also report [publicly] to the ICANN community. This is one way of doing. But we also have open RSSAC meetings. They are usually quite short and boring, but you're still very welcome to attend.

Next slide, please. We participate as one of the advisory committees. We have two members participating on the ICG panel for the IANA transition; or I should say properly, the NTIA stewardship transition for the IANA. We do participate to some extent in the accountability discussions, although as we've been thinking about that, we realize that RSSAC doesn't have much of a stake in those general multi-stakeholder discussions.

We also have liaisons to the ICANN Board and to the NomCom, so we have ICANN representatives on the board, one representative on the board and one in NomCom. These are non-voting members, but at least they can convey information both ways and inform and give advice as necessary.

Next slide, please. How's RSSAC constructed? It's composed of appointed members from the 12 organizations that operate root name servers. I said there were 13 addresses, so there's a mismatch there because one of the organizations actually operates two servers, two letters. So there are 12 organizations.

We also have liaisons, as I mentioned, to the board and NomCom, but we also have incoming liaisons from the three parties that are involved in generating the root zone. That means the IANA, the NTIA, the Department of Commerce U.S., and VeriSign, who does the actual editing and putting it on the servers. So we have them involved as well, and we also have a liaison to the Internet Architecture Board, which is on the technical side. The IETD and the IAB constitute the technical development of standards and the operations.

In addition to that, we also have something we call the RSSAC Caucus. This is a large pool of various types of experts and people who can contribute with their knowledge and experience because we're a very small group and we don't have all the experience and knowledge that we need to do the investigations and produce the documents that we see the need for. So we rely on more people to help us do the right thing, find what we need to look at, conduct investigations, and help us write the documents. We interact very closely with the caucus so that we get a good feeling and that we have the antennas out there so that we can get more input and also generate more output. So the caucus is very important to us.

Procedures. RSSAC, the formal committee, has a rather narrow or rather small task. Most of the interesting work is conducted in the caucus, but RSSAC is the administrative body that selects the work items so that we can prioritize and say, "This is the most important thing that we need to work on right now," or maybe a few ones in parallel. We appoint work parties in the caucus. We ask for someone to lead the work on creating a document or conducting an investigation. They will then invite other people to join them and report back when there are results to report back to the RSSAC. RSSAC will then take formal action, which is usually decide to publish the document or to give advice or make a statement or something like that.

The formal committee meets at ICANN meetings like this, where we'll have a continued meeting later today. We were in session all day yesterday. We also have regular telephone conferences, so every month we have telephone conferences where we forward our work and continue.

Next slide, please. Yes, we also do boring administrative stuff, like appoint liaisons and elect chairs and create the actual procedures and processes that we use when we work. So we have one document that describes how do we conduct our work, and that's our actually very first document, RSSAC 000.

RSSAC was a very small and slowly functioning body, and it didn't used to meet at ICANN meetings before because most of the people of the old RSSAC were engineers that didn't attend ICANN meetings because this is mostly a policy-related meeting. So RSSAC used to meet at the IETF meetings, which is the standardization body where the engineers typically go. So the engineers were at the IETF, so that's where the meetings were conducted.

Over the past three or four years, we have reshaped RSSAC into a new model, which is the one I described here. We've also realized that we need to be closer to the ICANN community, and therefore we have moved our meetings from the IETF to the ICANN meetings. We now meet regularly here.

As I said, we have to Co-Chairs. It's myself and Tripti Sinha, who unfortunately couldn't be here today. The caucus also creates transparency. The caucus is very open. Anyone who wants to join the caucus is welcome to apply at least, and we haven't turned anyone down yet. That means that since the work is created there and the mailing list is open for the caucus, that's where the transparency happens, where people can see what RSSAC works with and how it proceeds.

But it's actually transparency both ways because in order to join the caucus, you have to give a formal statement of interest. You have to say, "Why do I want to join the caucus, and with whom am I affiliated? Whom do I work for? Where's my very short version of my CV?"

On the documents that are published, we always list the names of the people involved in creating the document so someone who reads the document can see, "Oh, Lars-Johan Liman was there writing this document. Let's go to his statement of interest." They are published on the web. You can look them up on the RSSAC pages. "Oh, Lars-Johan Liman. He works for Netnod. Okay. I wonder how that influenced his work with this document." Then you can make your own judgment about that. That creates another building block in the transparency that we try to create.

Next slide, please. All members of RSSAC are also members of the caucus, but that's just a small part. The bigger part is people that have applied and people we have invited to the caucus.

Next slide, please. Yes. We talked about most of this. When the caucus receives a clear work statement, we always write a document and say, “This is what we’d like you to do this time.” We ask for a document leader to volunteer. There are the usual timelines and expected outputs and so on.

But another point here is the last bullet, which means that if you have a work party in the caucus who creates a document and there are people who really disagree with the content, that must be listed in the document. So if there is such disagreement, there will be a section in the document saying, “The following people disagreed with the following statement for the following reason.” It’s also a way to create transparency around the fact that people weren’t agreeing, that there was no unanimity – ah, difficult word. So even though we only have a rough consensus, we can still want to give advice, but then we also include the fact that other people disagree, and hopefully the reason for doing so.

The Caucus Membership Committee is a small subcommittee which only does the administrative tasks of receiving the applications, making sure that there is a statement of interest and so on, and then eventually forwarding the names to RSSAC for a formal approval. As I said, we haven’t denied anyone.

The Membership Committee is three people. It’s Kaveh Ranjbar, who works for the RIPE NCC and operating K Root. It’s Tripti, and it’s also Paul Vixie from C Root. They have a staggered rotating system, so we will eventually slowly rotate them.

Now the coming slides are usually presented by Tripti in her role in the Membership Committee. So we’ll just go through them quickly.

The Membership Committee maintains the continuous stream of our membership to the caucus and manages the membership. If someone wants to quit and leave, that’s of course not a problem. But we need to track of who is actually in the caucus. They also update the RSSAC, the formal committee, about what’s going on.

I think we’ve gone through most of this. The actual process for getting into the caucus is to send a message to RSSAC-Membership@ICANN.org. The Membership Committee will go through these statements. They will make sure that there’s a statement of interest. They will inform the candidates about the process and what’s going on. The Membership Committee meets every other month, so there’s two months between their telephone meetings. Then they work through the applications.

If someone hears about ongoing work in RSSAC and says, “Ooh, I want to contribute. I want to help with that specific issue,” that you are already working with, there is an expedited version so that we can get people in quicker than the two-month cycle if necessary.

If you want some more information, please, again, send messages to RSSAC membership. There is also a document which describes how we interact with the caucus and how the caucus works.

Moving from the caucus to what we're actually doing in RSSAC and have done recently, RSSAC has a series of documents. It's rather short so far because the old version meeting at the IETFs didn't produce many documents, and they were not numbered. So the process of creating documents wasn't very clear.

We did change that, so now we have a very clear process for producing documents. Apart from the 000 that I mentioned, which is internal, describing our procedures, we have one which is ready for publication but isn't published yet, called RSSAC 001, which is called Service Expectation.

Now, if I remember correctly, I'm actually going through these on the following slides. So RSSAC 002 is published. It's an advisory on measurements on the root server system, and RSSAC 003 is a report on the time-to-live values in the root zone and how they relate to DNSSEC signatures and so on.

We also have created various statements – for instance, on the signature validity period – and also statements on the accountability stuff that's going on.

Active work parties. We have one work party which is working right now. We realize that the names of the root name servers, as I mentioned, are denoted by letters. a.rootservers.net, these are domain names. You can look up the IP addresses for these using normal DNS.

The names have been in use for more than 20 years now. The change from the previous naming scheme, which was very non-deterministic – they just had ordinary computer names – was made in order to take advantage of the compression algorithm in the packets. We now have names which are not in the root zone. They have to go through .net and then into rootservers.net in order to reach the records that deliver the IP addresses for the root name servers.

One thing we're looking at is trying to rename the server so that they will instead live in the root zone and be contained in the root zone information. That, together with another several aspects of the naming, is being investigated by this work party. They have been tasked to come up with a recommendation whether to change the names or not. And if to change, what should we change it to?

We're also about to launch a small – probably not work party, but we found already in one of the documents that there was a technical error, where the document RSSAC 002 recommends a number of measurements that the root server operators should take. So they should measure the system and do it in the same way so that we can compare the numbers between the various operators. In one of the cases, it turns out that the numbers are skewed when you measure UDP packets, the single queries, or TCP packets, where you have a stream of information going through more packets. It's just a small technical detail, so we would probably not launch an entire work party around that. But we will talk to the caucus. We will ask for comments from the caucus for making this errata change to the document.

How are we doing on time? Okay. Okay. So RSSAC 001 is a document that's intended to describe the expectations that the Internet community should have on root servers. What should you expect from a root server? What do we intend to deliver to you? That document is going to be published in parallel, at the same time, as a document from the IETF side. It's published by the IAB, the Internet Architecture Board.

There are older documents that describe root servers and root server expectations. They were all in the RFC series from the IETF. But we realized that the newest one of those is more than ten years old, and things have changed on the network, as you may have noted. When we sat down and started to talk about writing a new revision of that document, we realized that, actually, it contains both requirements on the DNS protocol. So which parts of the DNS protocol should you expect that a root server can handle?

But it also had operational requirements saying things about capacity, about network connectivity. These are things that are not directly related to the quality of the protocol. So we designed it to divide the document in two separate ones, where the Internet Architecture Board is the right body to set the requirements for the protocol. They are responsible for, so to speak, the DNS protocol, so they can set the requirements for what should the root server do from a quality standpoint.

Then we had to find another body to deal with the operational side. Actually, there is none obvious. You could think of something like NANOG, the North American Network Operators Group, but again, they have counterparts in all parts of the world, so there's no single international body for that.

But then we saw of RSSAC and realized, "This is probably a good place where we can put operational requirements on the root servers and have them published." But we want these two documents to go out at the same time because they refer to each other and they are actually a pair of documents that are tied together, even though they are published by different bodies.

We're still waiting for that RFC. It used to be my fault. It's for the moment not my fault, and I hope to get it out really soon because it's actually written and we only need to have an okay from the few people that have asked for a few changes to it.

Next slide. So RSSAC 001 is the operational side, and it talks about the infrastructure around the root servers, the accuracy and availability of the servers, and capacity. Diversity is a very important thing because the root server operators all make their own decisions about how to provide the service, and that's our best strength because I decide, together with the Netnod staff, we may decide to operate in one way, and the Internet Systems Consortium, who operates F Root, may design something completely different. And again, University of Maryland, D Root, will take a totally different approach to how to do it.

The important thing is that you get the correct answer. The important thing is not how we do it to make sure that you get the correct answer.

By doing it differently, we are not vulnerable to a single type of attack. It we all used a special version of Linux, and that Linux version was attacked in this special way, had a vulnerability, we would all go away at the same time. We don't because we use one version of Linux and maybe they use Oracle Solaris on one site and they use IBM servers on another site. So that diversity is very important in the root server system, and we don't want to have a single organization take care of it because they will not create the same diversity.

The document also talks about monitoring and measurements and about how to be public about what we do. It turns out to be a rather problematic thing because reaching out to people when people don't know where to look for information is actually a hard problem.

RSSAC 002 ties into the RSSAC 001 because this document specifies the measurements that the root server operators are expected to do on their systems. It's simple things like counting the number of queries per second and looking at the number of clients. We don't care who they are; just how many clients do we have to our systems. This was all triggered by the New gTLD Program. When ICANN opened up for new top-level domains, we were starting to add top-level domains to the root zone in a much higher pace.

Before the gTLD Program changes the root zone, additions of new top-level domains happened, what, every two years? Three years? Now we have five per week or ten per week. The rate of change got a lot steeper. This is a new thing for the root servers and the root server operators. We want to monitor the systems carefully and to do so over a long time so we can see trends and we can see when things starts to happen. If there are problems with the system, we want to be able to compare to before now and see how we can fix the problems.

So far, we have seen no problem whatsoever. None. Nothing. We do see constant change with the number of queries, slowly increase, but that's the general trend on the Internet, and it's been like that for the past 25 years. So that's not a worry.

We also see certain attacks going on from time to time. That's also known use. That happened before the gTLD program and it happens after it, and there's no actual change to that patterns. I don't say that the attacks are not a problem because, of course, attacks are always a problem. But it's a problem that we can handle, at least as they're going on now. There's no news there. So RSSAC 002 is helping us to do this and compare the numbers.

Next slide, please. As I mentioned, a number of queries. Also some latency in the system. We received the zone file from VeriSign. We have to copy that to, in our case, Stockholm, Sweden. We have to copy that out to all our instances of I Root. There are more than 50 across the entire globe. We have servers in New Zealand, in Singapore, in Tokyo, in Beijing. We have servers in Mumbai. We have servers in Johannesburg, South

America, North America, and Europe. We have to copy that file from Stockholm to all these servers.

From the moment we receive a notification from VeriSign that there's a version of the zone file, in our case, until all the servers for I Root are updated, that's usually below seven seconds. That's one measurement that we have here. How long does it take to propagate the zone file to all the servers?

There are a number of other DNS measurements that we cannot keep track of for statistics and that we can go back and look at if we need to compare the numbers.

Next slide, please. This is just to give you a picture of what the statistics can look like. This is from the public page for K Root, operated from the RIPE NCC in Amsterdam. This is probably a bit old. Yeah, you see, it's from 2014. But it still looks roughly like this. These are their servers, and the colors are the various servers across the globe. They add them together so that you get the total number of queries.

You can see, at the bottom, the servers at AMS-IX in Amsterdam, DENIC in Frankfurt, I suppose, and LINX in London. NAP is probably – I'm guessing now – the NAP of the Americas in Miami. Those are their big servers. Also Tokyo. Then you have a lot of small servers on top of that.

There is more information to obtain about the root servers. We maintain a joint webpage. Now, this is not RSSAC. This is the root server operators because we try to maintain the difference. RSSAC is the advisory body to ICANN. The root server operators is the group of organizations that provide the service. For truly operational stuff, you have to talk to the root server operators.

To find them, you go to www.rootservers.org, as it says at the top right there. You will find a combined map, where all the organizations work together to indicate where we have placed our servers. You can see where they are placed across the entire globe. You will note that many cities have a couple of servers operated by different operators

I was going to finish there. Oh my. This is the fun part. Can someone please help me with – yes, please [inaudible].

UNIDENTIIFED MALE: Thank you. I'm [inaudible] from [inaudible], first timer in ICANN as a fellow. I'd like to know a few things. I saw RSSAC reports of root server TTLs. It is the root servers by literal agreement with any party to set any [inaudible] over there, like in a country. But you are actually giving reports to some organization about the availability of root servers around the world. Do you have any policy to give the root server organizers or maintainers that look your servers are real acquisitions and not available in this part of the area, so can you please [run or] something? Do you have any policy like that? I have another question after this.

LARS-JOHAN LIMAN: Okay. We'll take that one first. To begin with, there are different TTLs involved here. There is the time-to-live for the DNS records. That's what I was talking about before. But there's also the response time that you're probably thinking maybe of the TTL and the IP packets. That's a different type of TTL, and that can indicate the response time, how long does it take for the client to get a response from the root server. That's an important property.

But the root server operators, the 12 organizations, deploy servers according to their own policies. There's not a common policy for all root server operators. So each and every one of them have different policies, and the only one I can speak for is for Netnod.

In our case, we definitely try to find the information about where it's problematic to reach our servers. When we do that, we try to take into account whether it's possible to reach another operator's servers.

I'm taking a country out of the blue: Uganda. If it's difficult to reach our server in Uganda, but Uganda has F Root, B Root, and

D Root, then it doesn't make sense for us to place a server in Uganda because they're already taken care of by the others. That's another thing we look at.

We try to identify areas where it's difficult to reach any root server. But that's not the only thing that influences our selection because there are other things that are very important. Is there a good place to put a root server? We can easily detect that it's very difficult to reach any root server from the Antarctic, for the South Pole. But it doesn't really make sense to put a root server on the South Pole because there's not many users there. There's probably not a data center where we can put the servers. There's no one there that can help us manage the computer when we need to have it reset or change the hard drive. And there's no one who's willing to pay for it.

There are many other things that influence the selection of these Anycast sites, and every root server operator has their own policy for how to do that.

I would argue that most of the operators very much welcome initiatives from communities to deploy root servers. I certainly know that we do. If you feel that the root service doesn't work well for you in your region, please come and talk to me. I will not promise that we actually do place a server near you, but I very much would like to know that that's the case so that we can put it on the list of sites that we take into account when we deploy.

UNIDENTIFIED MALE: Hello. I am [inaudible] from Argentina. My question is about what is the difference between root servers and mirror root servers? Because as far as I know, they are [inaudible] original or main root servers, [inaudible] mirror root servers.

LARS-JOHAN LIMAN: Let me now kill a myth. There are no mirror root servers. Please do not use that term. They do not exist. There are root servers. They are all alike. They all function in exactly the same way.

The thing that may confuse you into believing servers and mirrors is that you say there's only 13 IP addresses talking IPv4, and that is correct. But we cheat. We play tricks on you. The trick we play is that we use a technology called Anycast. I Root is only one IPv4 address. It used to be, if you go back 15 years in time, we had one server in Stockholm on that specific IP address. That was before we were using Anycast.

Today we have more than 50 servers with the same IP address. Our server in Perth in Australia has the same IP address as our server in Japan, and the same as the server San Francisco, and the same as the server in Miami. They all have the same IP address.

When you talk to I Root by using that IP address, you get to the nearest one, the nearest copy. They all are identical. They all have the same information, the same IP address, the same everything. The routing system that carries the packets on the network will forward the packet to the nearest one.

Now, "nearest" is kind of a strange denotation here because the routers decide what "near" is. But it's well-defined. In the routing system, there is an algorithm that says, "Send it this way."

UNIDENTIFIED MALE: And it jumps, yeah.

LARS-JOHAN LIMAN: Yeah. And that's taken care of by the Border Gateway Protocol, the BGP system, for sending packets. There's no difference to describe because they are all identical. All these servers have the same IP address, so how can we from Stockholm reach the one in Perth? It should just go to the Stockholm on in our case. Every of these server machines has two addresses. It has one which is the same on every place, and then it has one in the rear, which is unique for every site. That unique address is what we use to update the server with the correct information. Then we can reach them all because they have the unique addresses.

So, please. All are the same for I Root. All are the same for D Root. All are the same for F Root.

UNIDENTIFIED FEMALE: [inaudible]

LARS-JOHAN LIMAN: I'm happy to talk afterwards.

MANUEL HACES AVINA: Okay. My name is Manuel from NIC Mexico. I was wondering, why has it ever been 13 root servers? Is that like a convention, or is that like a history? Could you elaborate there?

LARS-JOHAN LIMAN: Actually, technology. If you go back in time, you have to remember that the root server system is very old. It started to be deployed in the 1980s, around 1982 or '83 or '84. That's when the DNS standard was created, the protocol specification.

Now, the Internet 30 years ago was something completely different from what we have today. One of the things they worried about was the size of the packets. The old DNS standard, which is still the standard in use, specifies that the packets may not be bigger than 512 bytes, 512 characters.

The most important DNS queries you can send is to ask for the list of root servers: “What is the current list of root servers?” In order to fit in that 512 bytes, there’s not room for more than 13. That was actually the reason for renaming from the old names to the rootservers.net because by taking advantage of the compression, we could make for room for a few more. We went from 9 to 13 by just renaming them. So that’s the old limitation.

Now, you can argue that, on the modern Internet, we have new extensions to the DNS protocol – the eDNS (extended DNS) – which most servers used, where you can use larger packets.

The limitation? We want to be very, very, very secure that everyone can reach the information, so there has been no decision to extend that beyond the current system. Another reason is that the obvious process for extending was Mr. Jon Postel. He unfortunately passed away in 1998. Currently there is no process for changing the root server system.

Those who operated the servers in 1998 are still operating the servers now, and what we do is start to use Anycast to try to cover as large parts of the globe as we can.

UNIDENTIFIED MALE: [inaudible] from India. As you mentioned, as of now there is no technology limitations that the number of root server instances or the letters can be increased. Is there any proposal to be submitted by RSSAC, or is any discussion going on, especially in the content of the IANA stewardship transition?

Second question –

UNIDENTIFIED FEMALE: [inaudible] just one.

LARS-JOHAN LIMAN: I am aware of a lot of pressure for creating more servers. I will phrase it like this. If more root servers operators should be added, it should be done so for the correct reasons. It should be done so because there are problems to solve, and the problems should be technical. If adding another organization operating root servers solves technical problems, then we should look into it.

JANIC DOUMA LANGE: I hate to say it's over because I know you have a lot of questions. All I can say is, if you're free – I know you have a 10:00 – it's up to you. Then you can follow outside so that they can prepare for the next session and ask away if you are free.

LARS-JOHAN LIMAN: Yes, I am free for another, say, half hour, because RSSAC is meeting with the board later on here. I'm happy to take questions. Also, in the presentation that you will send by e-mail to all of you, my e-mail address is in there I believe. Otherwise, Yannis has it. I'm happy to take questions over e-mail.

Please stop me in the hallway if you see me. I like to interact. This is actually the best session in all of ICANN because you have so many good questions. The others don't ask. They just plow on, but you want to know. That's good.

JANIC DOUMA LANGE: Well, I'm going to say thank you very much first. Anyone who would like to follow Liman out, you are welcome to because he just offered you another half an hour of good discussion. There's some tables outside, so I know you can find a place to set up so that they can get prepared here for the next session.

There is a sign-up sheet going around, so before you exit stage left with Liman, please make sure you've done the sign-up sheet. Liffey B at 5:30 for the Nominating Committee and some other speakers: Security and somebody else at 6:30. Can't remember who.

[Martin], quickly. You wanted to say?

[MARTIN]: Yes. NPOC is organizing informal outbreaks, informal talks, today and tomorrow. Both days it's on the fifth floor in the hall. The topic for today is we are trying to create a clearinghouse for NGOs so NGOs can have legal consulting for free or at a minimum cost, basically to help them on how to protect their name on the Internet, or how to protect from abuse – those sort of things. Tomorrow, if DNS has anything to offer to the migrants and refugees. So in case you want, it's 12:30 to 1:20. It's completely informal. It's actually something we're trying. Whoever wants to come, you're welcome.

JANICE DOUMA LANGE: Great. Once again, thank you.

LARS-JOHAN LIMAN: Thank you all.

JANICE DOUMA LANGE: Again, you can follow out as long as you've just signed the sign-up sheet. I'll see you at 5:30. Thank you.

LARS-JOHAN LIMAN: Thank you.

[END OF TRANSCRIPTION]