MONTREAL – Tech Day - AM Session
Monday, November 4, 2019 – 10:30 to 12:00 EDT
ICANN66 | Montréal, Canada

EBERHARD LISSE: Good morning, everybody. This is the 40th Tech Day. My name is Eberhard Lisse. I'm the ccTLD manager of DotNA and chair the working group that organizes it. It's quite a good problem to have that we seem to have standing room only. So, probably in the future, we'll have to kick out the people from the opening ceremony a bit earlier and use their room. A nice problem to have.

As usual, I'll go through a quick opening, review the agenda quickly, and then we'll kick off the proceedings. We will start with a presentation from Benin. Yazid is a fellow and they have done some sort of a hackathon. And when I heard the two words Raspberry PI, I thought that's something that we may want to hear about because I always like to have to hear about cool little technical gadgets and things that can be done with them. And in resourceful locations, it's always important to have ways out.

Then, we will hear about the Canadian Internet Exchange Point Landscape from Jacques Latour. Then we have our SSAC presentation. Andrei Kolesnikov will speak about application port scanning. That's an issue that they have identified that we are going to hear about. Ondrej Filip from DotCZ is a double threat today. He will make a presentation about DDoS prevention at speed and he will wrap up the proceedings. We'll come to that.

Lunch is served catering, unfortunately. Then we will hear about RDAP deployment and Mark Svancarek from Microsoft will tell us what they do with WHOIS at Microsoft. He gave an excellent presentation in Japan about deployment of IPv6 on [inaudible] so I'm looking forward to that.

We will then have two machine learning presentations from Qatar and from Japan, or the respective NICs. I put those two presenters in contact with each other so that they can sort of avoid duplication and make things fit better to each other.

We then have the host presentation since the cyber justice faculty of the local university is doing this. I can imagine that there may be also some machine learning issues in their field of expertise, so I put them together and also made them contact each other so that maybe they can get some synergy out of this.

Bruce Tonkin will then speak about penetration testing AUDA which is something with regards to security that I found very interesting, especially larger registries with more stuff are more prone just of a number of stuff that can be [inaudible] in various more places where there's one or two people and somebody gets an email that one doesn't really think is correct is probably easier to manage but on larger companies, like AUDA, it may be more difficult. But I'm quite interested in hearing what the outcomes are.

Then we have so many presentations that I have two slides for the afternoon. We'll then hear from ICANN about the root server early morning system that they are working on.

Martin Wullink will give an Entrada update. Then we will hear about ICANN security practices from ICANN Org again. Then we will hear a presentation from the Indian NIC. They recently migrated from Afilias to Neustar if I'm not mistaken and they will share their experience with it.

I'm quite interested personally on the communication plan they used to inform a large number of registries in a huge country with different languages because I am vice chairing the retirement policy development working group where our policy will involve that, if this happens, a plan would be made and it should include also communications plan. So, how they have handled their situation might inform our working group, so I pointed this out to the presenter that he should spend a little bit of time on it, and if he doesn't, I will hold him to it.

Then, a late edition is quantum DNSSEC. We have recently heard that Google had a proof of concept. Of course, they make a little bit marketing about it and IBM, as the competitor, was not impressed and so on. But quantum computers seem to be coming and we all are aware, to some or other degree. I don't understand much of it but the algorithms that we use currently may not withstand [inaudible] challenge of quantum computing. So, it is incumbent on the designers of DNSSEC or whoever writes the upcoming RFCs on it to look early. Do we need to change our algorithms? What algorithms are we going to use that could withstand a quantum challenge?

Then, Ondrej Filip is going to do his second job. He was volunteered for his other vice chair to do this. He will wrap up, taking notes during the day, wrap up. [inaudible] and then he will give me his notes and I will shamelessly plagiarize them into a report that we then publish for public consumption.

Without further ado, Yazid. Sorry, I must confess I have a bit of a problem with the surname. Yazid Akanho is a fellow. I always like to communicate with the organizer of the fellowship or the local contact, the ICANN contact, because I really want fellows who are starting out on this to have a venue that they feel where they can come to and talk to friendly, and now even large audience, about topics that they find interesting and often it turns out that we find them interesting, too.

YAZID AKANHO: Thank you, again. Hi, everyone. It's a big pleasure to be in front of you here. I'm Yazid Akanho from Benin in west Africa. I'm a fellow, like Lisse mentioned that. This is my first physical attendance to ICANN meetings. So, it's a real pleasure for me. I'm greeting also all my friends, fellows also, and [inaudible] who is my mentor during this fellowship program.

So, I'm here to share the experience we have been in during our Benin DNS Forum. We mainly talk about the hackathon which is one of the activities of the Benin DNS Forum. So, we basically used Raspberry Pi to reproduce the DNS architecture. So, we go through the slides and I will explain.

So, basically, what is Benin DNS Forum? Benin DNS Forum is a local forum dedicated to domain name system and also Internet. It was actually inspired from Africa DNS Forum. We started in 2014. We started Benin DNS Forum in the next year, 2015.

The idea was to build a local ecosystem on Internet and DNS issues but also to improve the engagement, the gender and diversity also on Internet governance areas. Of course, also, to improve the technical capacity and entrepreneurship locally in Benin.

Actually, we have five local associations which are working together to make Benin DNS Forum a reality. So, it is a yearly event. It happens once a year.

Through the five years of Benin DNS Forum, we tried too improve and to innovate also. So, we basically started from one activity at the first year and now we have reached six activities quickly. Woman DNS Academy, which is dedicated to women, empowering in technical capacity. We have a multi-stakeholder meeting with managers to explain them what are the best practices they must follow in terms of DNS. We have a hackathon which I'm talking about here, actually. We have two public forums open to the public where we explain what is Internet, what is DNS, how it works, why we must care about security also when we are online.

And this year we introduce Internet City to help the enterprise and also the ICT managers to show what they have as product. You can go to our website, dnsforum.bj to see what we are doing and we have many reports also. We have [inaudible] since the first edition.

You can see the graph here, the participation – the growth of participation. So, the first year we started with almost 100 people but now we have reached 700 people, mainly because we tried to diversify the activities.

So, DNSathon is a hackathon on DNS. It was introduced in 2017 at the third edition of Benin DNS Forum because the local community wanted to understand how DNS worked. Even [inaudible] they wanted to really understand how DNS worked. So, the organizing committee tried to think about hackathon which will be fully dedicated to DNS.

Every year, we try to receive 50-60 participants. We make a call. They submit and we select some of them based on diversity and profile also. So, they work on a really collaborative approach to contribute and to implement the prototype. They actually work on implementing the root server, the registries, registrars, and also resolvers to understand how it works.

Initially, we were also wishing to implement DNSSEC but after two editions, we were not able actually to do it. And we use Raspberry Pi, like I explained why, because it is quite cheap and also to make really the attendees understand what they are working on so they can see it actually and they can directly work on it. Yes. we also prepared a white book to help other communities also to reproduce the same.

Basically, to build this DNS architecture, we separate the participants into six groups. One worked on the infrastructure building. One worked on building the root server. Other work on the registries. Other group work on the registrars and resolvers. And the sixth group worked on the

documentation, so they collect all the documents from each group to produce, to build one common document.

We have one coach, who is Alfred Arouna. I think most of you, or some of you, may know Alfred Arouna. He's really also engaged. We have [Romano] [inaudible] and [Mattais] and myself also as assistant.

So, we start by brainstorming session to help everybody understand at least to know what we want to do. And on the goal, we have debriefing session to help each team to show what they are doing, what are the challenge they are facing and how they want to overcome it.

Also, during the hackathon, we used a lot of different solutions to help also the participant understand that DNS ecosystem is based on different software solution. This is how DNS works, actually. And also it is opportunity for them to understand how Bind can be configured, how we install and configure [Power DNS], Unbound, etc.

This is the architecture we built. It is quite simple. We have a router, two switches. We separate the ports into VLANs. We put each Raspberry Pi into separate VLAN to also make this architecture not one common architecture.

And in [inaudible] we have all the attendees who are connected to this switch and they can remotely connect and configure the Raspberry Pis. So this is how the teams are composed and each team work on their stuff.

So, in terms of outcomes, the hackathon helped the attendees to better understand how the DNS infrastructure works. Root servers are

[inaudible]. Nobody can really access root servers and type comments to see how the system works and seeing the Raspberry Pis in front of them, accessing them directly, configured them. Make even them more confident in [inaudible].

The system was fully working. All the components worked. We have actually all the results which are consolidated on [a public share] on GitHub. So, the router configuration, all the switches configuration, all the Raspberry Pi configuration have been published on this GitHub, so we can … At the end of the presentation, we quickly go on this. You can see the infrastructure. So, anyone can fully rebuild this infrastructure as soon as you have the number of Raspberry Pis. We use seven to eight Raspberry Pis to build it.

We are currently finalizing the white book. The white book is actually based on the GitHub but we put some text to help understand the step-by-step system.

What is next? We have three challenges. The first one is to implement IPv6 over this IPv4 infrastructure to also help the participant understand IPv6 and also understand [geo stack] concept. DNSSEC, yes, we would like to also deploy DNSSEC in this infrastructure. It will be quite difficult but we can do it simple. Thank you very much.

Also, those new concept of DNS over TLS and DNS over HTTPS also want the attendees to understand why those concepts are new and what they bring [inaudible].

We are available to support any community to reproduce this. We really think that it will help our communities, mainly in Africa and anywhere else in the world to understand what is DNS, how the complete architecture works. We have the motivation to also learn the new stuff I have put over there, and yes, we have the engagement since more than five years now.

Yes, those are some pictures to show how the participants work and how they challenge themselves, how also they present what they have done and the challenge they are facing and how they want to overcome them.

We tried to make them find a solution to the problems. We can guide them, but at the end, they are the ones who provide the solution.

At your left, this is 2017 edition, and at your right, this is 2018. So, you have the root, you have the TLDs. And the root, we build our own root. So, we even [inaudible] the root [hints], the content of the root [hint]. We removed the [13] root servers and we put our own root server IP addresses. So, it is a local DNS infrastructure completely. You can go to the GitHub. You will see all the details there.

Thank you very much. Those are the teams in 2017. It is also time for me to thank those who trust on us since the first edition of Benin DNS Forum. It was not easy. We have ICANN. We have Afilias. We have ISOC and we have a lot of local support also. Since three years now, they have understood what is DNS, why we must care about DNS also.

If you allow me, Lisse, just quickly to go on the GitHub. Yes, this is the white book we are building. I hope by end of December we will complete it and publish it on our website, dnsforum.bj. So, this is the GitHub. Please help me click on 2018. Yes. Thank you.

So, everything is there. You just have to click on it. You will have everything. Even we put the monitoring in place, so all the [JSON] dashboard there, you can go through it, put your comments. You can get in touch with me or [inaudible], anyway you find us to give your comments. Thank you very much.

EBERHARD LISSE: Two things. I think we made a mistake. We should have put that at the end as the highlight presentation because that was very impressive. DNSSEC is very simple. It's very easy to do. It's just expensive if you're doing it in hardware but you can do it [inaudible] with Bind. We can get in touch and we can show you how [inaudible]. It's simple to do for teaching purposes. It's simple to do. It's just not [auditable]. Any questions?

One more thing. All the presentations will be posted. Any links in the presentations that are clickable will be accessible so you don't need to make photos of the webpages. You don't need to write the links down. All these presentations will be published, so you read the presentation, click on the list, and you reach the GitHub. You must come to the microphone, please.

One thing. If there is remote questions, remote questions have got precedence and I will allow about four minutes worth of questions. The lady was first.

ROLLA HAMZA: Good morning. My name is Rolla Hamza. I'm an ICANN 66 fellow. Actually, I have four questions, but for the time limits, I'll do two only.

EBERHARD LISSE: One.

ROLLA HAMZA: There is a gap between participation of women and the man, women and men. So, how do you intend to close the gap between the participation to increase women participation? This is first question.

EBERHARD LISSE: Sorry. Just one question, please. Just one question. We have got three people. We have got four minutes.

ROLL HAMZA: Okay. That's all. Thank you.

YAZID AKANHO: Thank you very much. Yes, you are completely right, and this is even why the second edition went through Woman DNS Academy, to help women understand how the ecosystem works, to help them get more

engaged in the ecosystem, help them get more confidence to configure the systems.

Also, to select the participant. Definitely, when you have maybe 100 who applied, you only get maybe 10 ladies. You can't go more than those who applied and we are trying to change this balance. Thank you.

EBERHARD LISSE: Being from an African country myself, I feel in part to also mention something. We don't really need to push woman participation for the sake of women participation. We need to remove barriers. If no women take up the business, they will not come to this. We should be as barrier-free as possible so that we make it easier. [inaudible], I think, in my view, is very good.

I'm taking [Yoshi] and then the question from here. You sat down again. Okay. Then just [Yoshi].

[YOSHI]: My name is [Yoshi] [inaudible] from JPRS. My question is the participants on the picture are very young, so I was very surprised. So, I'm wondering where their affiliations? Did they come from universities or ISPs? That's my question.

YAZHID AKANHO: Thank you for your remark. Yes, you are right. We also tried to balance. This is one of our objectives. It is true. I cannot hide this. In our universities, we just learn theory. I was one of them. I know what I'm

talking about and we are trying to overcome this. So, we tried to select 70% young people coming from universities, third or fourth year of ICT studies exactly, and 30% are engineers who are already working. That is usually how we try to balance. Thank you.

EBERHARD LISSE: Capacity building is extremely important in developing countries, as we all know. This is a young man's game. I've been doing this for 30 years now and I feel old when is see what young [inaudible] come up with nowadays. Thank you very much. This was one of the more impressive presentations we had for quite a while.

All the presenters have got their email in the agenda as links, so you want to ask further questions, just email him directly. If it is interesting, the outcomes, you can always send it to me and we publish it on the list.

The next one will be Jacques Latour from CIRA about the Canadian ISP landscape.

JACQUES LATOUR: Good morning. Welcome to Montreal. So, today I'm going to talk about the Canadian ISP landscape. Over the last ten years, we've done a lot of work but we haven't talked about it a lot, so I thought it would be a good opportunity to explain what a ccTLD can do to help its country have more ISPs.

But before, we need to talk about poutine!

UNIDENTIFIED MALE: I had it already.

JACQUES LATOUR: Did you like it? Yeah?

UNIDENTIFIED MALE: Yes.

JACQUES LATOUR: Very acceptable. So, if you Google Montreal, there are poutine tours. That means you can go find a nine poutine course. So, you go to nine different places. You taste nine kinds of different poutine. I suggest you don't have breakfast before and maybe buy some Pepto Bismol after. It's the place to taste them. And smoked meat.

A little bit about CIRA. So, we got 2.8 million domain names. We have four Anycast clouds. We have cybersecurity services. We have a booth over there. So, the money that comes in helps us invest in the community initiative and part of that is we donated $5 million in the last five years, about 130 projects. And some of these are IXP related. Also, over the last ten years, we've heled eleven IXPs in Canada to get built, to be put in production. So, that's what I'm going to talk mostly about today.

So, CIRA, we have three [ASN]. So, one for the registry. We have two clouds, one [AS] each. We have about 5,000 peering relationships

worldwide, which is pretty good. We have 2,000 unique peers. So, there's many peers that are repeating themselves at [inaudible]. We buy a lot of transit. We peer a lot.

But sometimes when we peer with large ISPs, they say, "We don't want to peer with you because you don't have enough traffic to peer with us." And that's a problem.

So, some of them, you want to peer your ccTLD directly with large telcos or ISPs and that's one of our key criteria in making our global reachability. So, I'm not going to tell which one they are. But some are in Canada. I'll go more in detail on that.

So, we built an Anycast that the overall goal of our DNS infrastructure for DotCA is to serve Canadians. So, we built an infrastructure global, to have a global presence, like most of us Anycast provider we have. And then, in Canada, we're trying to have as much as possible a local Anycast infrastructure on different clouds.

So, the goal is, if there is a DDoS, if there is an attack, Canadians are more resilient on this infrastructure and that's our number one objective. So, being able to be as resilient as possible in Canada to have as many ISPs as possible, to have DotCA everywhere, to have roots everywhere, and make our infrastructure better for Canadians. So, that's that ultimate goal of this, our Anycast network.

One of the visions that we have is to serve DotCA from Canada to all Canadians. That's our goal. So, anybody in Canada that types a domain name, you want to resolve directly in Canada to one of our name

servers. That's our goal. But we're not. And that's the biggest challenge that we have.

So, the biggest Canadian DNS site that we have is in Ashburn. That's where most of the Canadian traffic is resolved. And even in Canada today, we have large telcos, blue, that don't peer with – if you look at their logo, it's blue, so you can figure it out. They don't peer with us in Canada. And that's a challenge because that telco doesn't peer with pretty much anybody in Canada.

So, it's hard for CIRA when our strategy is to have infrastructure in IXP and peer with Canadians to keep all of our traffic local. So, we're working hard with different partners in Canada to fix this, but even here we're having issues in our own backyard getting our traffic local.

For some telcos in Canada, they don't even peer with us at the larger ISPs in the US because we don't have enough traffic to justify their policy, so sometime we need to fix that. So, the more we talk about it, the more people hear about it, and we can change the way people behave. So, blue log, if you're here …

So, over the last ten years, we went on a journey to grow, to add more ISPs across Canada. And I want to talk about that. I'll be honest. When I started at CIRA, I didn't know much about IXPs, so we had Bill Woodcock and PCH come in and do a report on the state of the Canadian Internet. That was a pretty good report. It enlightened us on what we should do or where we're at.

So, back then, we had two IXPs. One in Ottawa, one in Toronto. The one in Ottawa, it was not growing. TORIX was working well. So, what Bill said is, "What you should do is you should look at building eight priority IXPs in Canada." Remember, this is in 2010 or 2011 when we had this done. So, from Halifax to Vancouver to build Internet exchange point. So, that was a strategy, basically. So, that was a challenge to figure out how to do that.

The next priority was, "Well, you need to build 15 more in more cities. That would be the plan. Then, over time, the Internet will grow and then you're going to have requirement to build in those cities." So, this is report from PCH.

So, overall, they said we should have 22 IXPs in Canada in a major center and [try to build]. So, in 2012, that's where we had in Canada. We had TORIX and OTTIX. There were many other IXPs, but most of them were controlled by their research and education network and the peering policy was only for the education sector, not for anybody to peer. So, they were not real open public IXPs. They were closed. So, we tried to change that.

So, in 2012, we told the world – the told Canada – that we would start helping building IXPs. Obviously, everybody misunderstood what we meant because everybody assumed we're going to take over all the IXPs and Canada, CIRA is going to control everything and there was a lot of fake news, because it existed back then.

But the goal was to help the communities get together and to start … So, in 2013, we [inaudible] three new IXPs in Calgary, Winnipeg, and in

Montreal to form. There was a lot of challenges and issues around the perception of CIRA, that they thought we were going to come in and take over. So, this was working with them. We had to learn about IXPs along with the community.

And then, 2015, we added three more. Most of those started with local people in the community. ISPs, they got together. They said, "We need to start an IX." Some of them asked us to do a town hall and present and explain what it was. Others, simply like in Vancouver, there were multiple parties and we had friction with many of them. But in the end, the community got together in Vancouver and they actually built an IX.

Everybody told us in the beginning it would fail because it was only three milliseconds away from Seattle. Nobody is going to invest to peer in VANIX when Seattle is so close. And then something happened in the US that made a few people think it's better to peer in Canada. I don't know what that could have been. But VANIX is the second-fastest growing IX in Canada, if not the fastest. So, they have over 50 members peering in Vancouver.

So, at this time, people started to … We had a vision for all [AS] in Canada, all networks in Canada, to peer locally and exchange traffic together. So, we got a lot of momentum around this. At this point, we were starting to look at the government of Canada and large networks to peer locally here and exchange traffic in Canada instead of the US, because at this point, the vast majority of the traffic flows in New York, Chicago, Ashburn, Seattle. That's where Canadian traffic is exchanged

mostly today. So, we need to build infrastructure, get Canadians to trust it and then use it.

2017, we added [inaudible] IX. This is super telco controlled province with SaskTel. And we managed to build an IX. Now it's expanding to its second location within the city of Saskatoon. Small ISPs are starting to connect wireless ISPs. If you build it, they will come, and I think, for IXPs, it's true. It takes time but it happens. Then we added a few more in eastern Canada. There's a lot in eastern Canada. There's one more that came in.

OTTIX is an interesting use case. So, the hardest place in Canada to build an IX still is in Ottawa and today we have one but it took a very long time. The Internet Exchange was started in one data center. Then it got acquired by the telco. So it was not vendor neutral anymore. Then they moved it to another data center. It got acquired by a different telco, so not vendor neutral. So, they had to move the IX again three times and then on the fourth time, people gave up when it got acquired by Rogers.

So, we still, today, we're looking in Ottawa for a carrier hotel where people can break out their network. It doesn't exist. Everything is super controlled by the incumbents. So, there's a lot of work we need to do there.

So, with all of this, we have discussion … So, OGIX is under development right now. It's live. So, I should have put 2019 year and [inaudible]. Price Edward Island is an island and having an IX is pretty useful to be more resilient because when the fiber gets cut on the bridge

by a nice bird or something, they completely lose the Internet, so there's no DNS on the island. There's no … They lose name resolution. They can't do anything. So, having a local IX with some critical service would help.

Then, we did a couple of blogs on ArcticIX. So, in [inaudible] up north … So, that's the one on top. You don't think they have [inaudible] networks. So, all of the networks have a [inaudible] link and none of the networks on the island in this community, there is no inter-connection of networks. So, the kids at school, when they have to submit something, they've got to go down to Saskatoon and then the networks don't even peer there. They got to go in the US, exchange traffic to the other network, back to Saskatoon, up in the sky to the school network. It's crazy. The latency for DotCA when I got there was like 15 seconds to resolve a domain name. I don't know how they can work there. So, I guess there is good and bad time there.

So, building an IX, interconnecting local networks, getting local content available to the local community was an objective. And we're still looking for a facility but that's the biggest challenge we have there. So, there's nothing there, right? So it's pretty …

So, we're talking about adding more IXPs in Whitehorse, in Yellowknife, in Regina. So, it's not in the same order that PCH and Bill Woodcock predicted. The order is based on the community getting together, seeing a need, communicating [inaudible] the requirement that they want to peer, CIRA and other communities, we get together and we make it happen.

So, the challenges that we have. A lot of the Canadian traffic today still is exchanged in the US and we need to bring it back in Canada, so more peering. We need to bring more content provider in the IX. So, when there's a critical mass of ISPs, then the content will show up. We're actively working with government of Canada and the provincial to peer. Then, once you build that, you build an ecosystem and there's more transit, there's more value in the IX.

So, as a ccTLD, I think that there's a list of … So, we're vendor neutral. We don't start CIRA. We don't start new IXPs. So, we support the development of new IXPs. We don't operate them. The community, they support them. CIRA does not own any IXPs. They're all not-for-profit. They're all vendor neutral. They're all member based. The IXPs govern themselves. We don't govern them. And we definitely don't rule any of them. That's not our game.

So, we're there to support them and I think the strategy over the last ten years is working with more IXPs, more traffic getting routed locally. But we have a lot of work still to make Canada more autonomous infrastructure.

So, we started to build CA-IX which the Association of Canadian IXPs. There's a meeting today in Toronto. They're all meeting there, most of them.

And that's it.

EBERHARD LISSE: Thank you very much. I have one question. I'm not one for regulator or regulations. I know from my profession that the [inaudible] model is the same in the rural and urban areas because they fly all the women when they're pregnant to the bigger places. Is there not a place for a regulator to say that, especially in these remote areas, that the [inaudible] must peer? Can't it be enforced by regulation of legislation? Have you looked at that?

JACQUES LATOUR: Oh yeah. Everybody talks about getting the [inaudible], the regulator, to do stuff. But they have no say in this matter and I would think we want them to have any say in those matters either. I think it's up to the community to get together and raise the issue. I think CIRA is doing that. We're raising the issue. But it's all about money in those locations. They make a lot of money sending bits in the sky up and down. Then if you exchange locally, somebody is going to lose money somewhere. I wish we could but, otherwise, [inaudible].

EBERHARD LISSE: Any questions from the floor? I can take two. Can you identify yourself for the remote audience, please, and for the tape?

UNIDENTIFIED MALE: [inaudible] DotLB. I was not sure if you said that the IX's are interconnected or not interconnected.

JACQUES LATOUR: No, they're not interconnected. They're all isolated.

UNIDENTIFIED MALE: All isolated. Thank you.

JACQUES LATOUR: In eastern Canada, we have four IX that are really close to another and we're looking at building some shared services among them but not interconnecting them because it would compete with the local ISPs or telco.

PABLO RODRIGUIZ: Good morning. Pablo Rodriguez from the registry DotPR in Puerto Rico. How do you get to manage competing actors such as ISPs and others to trust each other to participate in the same place? It has been our experience in Puerto Rico, they don't trust each other and consequently are not willing to participate to connect to the one particular place.

JACQUES LATOUR: Well, that's when you build a community. If there is value in the IX for the peers to connect, then they're going to get together and connect. In some of the cities, we had competitive ISPs that didn't trust each other and we ended up starting with two Internet exchange points.

So, one was operated, one was prime with one ISP. The other one … So, we had two in the community, but over time, you work on getting them to trust and you merge them together. So, we did that a few times. I

think we still have one city that has two IXPs but sometimes competition is good.

EBERHARD LISSE: The problem that Pablo has mentioned is the same everywhere. We have the same thing. [inaudible] was also a big drama to get them to talk to each other. And we only have got one IXP in the middle of the country but we have a country that is smaller than Montreal citizen-wise so it is not a problem.

JACQUES LATOUR: I just have one more comment. In Canada, there is one large incumbent. Telis. They're the only one that peer with CIRA in Canada and they connect to the IX and they only peer with us to resolve DotCA local. It's a first step and hopefully not the last step from these guys. Thank you.

EBERHARD LISSE: Okay, thank you very much. Next one is Andrei Kolesnikov.

ANDREI KOLESKNIKOV: Good morning. My name is Andrei Kolesnikov. I am part of the SSAC, Security Stability Advisory Committee, of ICANN. My presentation today is a little bit out of scope of the traditional DNS daily work but it's very interesting and very intriguing.

First of all, let me say a few words about SSAC. We are 38 members appointed by the Board and the role of SSAC is advice to ICANN

community and Board on matters relating to the security and integrity of the Internet naming and address location system.

Since 2002, we've got about 106 publications and our experience within following items. Addresses and routing, DNS and DNSSEC, registry and registrar operations, ISP and networks operations, DNS abuse and cybercrime, internationalization of domain names, and ICANN policy and operations. Also, as far as I understand it, as far as I remember, for the Tech Day, traditionally there is very interesting technical agenda. Thank you. It's always very interesting to be here and hear about the news and interesting problems and interesting solutions. So, let's move forward.

Here is [inaudible]. It came out from Facebook. I was following my friend who is a technical guy. He's actually writing the code. And he posted on Facebook the screenshot of his browser while doing banking. I mean, the guy was in a close relation with his bank transferring money doing whatever, and somehow he turned on the java script console on the web browser and found that the script – the bank, basically – is scanning his local computer.

Then, we found that it's not only banks. There are many sites highly loaded, very popular – I'm at least talking about Russian sites. It's not extended to the rest of the world but we'll figure it out later. They do scan the local host through the java application. They do not give notice to the user. There was no warning, nor little hidden text in the banking license agreement. No consent, of course, asked.

So, what it does, basically, is java script runs in the web browser on the user computer and generates an HTTPS request to the local host, like this you see on the screen – 127001. And the port number 6980. But it's a lot of ports being scanned.

Also, when you start digging down, you see that java script is obfuscated to prevent routine inspection. Basically, you do not see the source code of the script. You may only recover some of the functions it does.

And, of course, it's not yet clear what the purpose of scanning is. It looks like this. If you open your browser, you can turn on the java console for the windows. It's ctrl+shift+J and for the Mac it's command+shift+J. Basically, you see on the screen, it's scan of the different ports. The console shows you what's going on, basically.

Also, the same story we found out with ISPs, one of the largest ISPs in Russia. The attempt to recover the script shows you the script is obfuscated. You can only make an assumption that what it does, it measures the delays between requests to the local port and the reaction of your TCPIP stack, basically, on your personal computer. On the right side, you see the number of the ports. It's basically all kinds of different ports trying to see if [inaudible] open or not.

There's some technical things but I'll go fast through it. Basically, you can reproduce this behavior with HTTP [inaudible] to the local number, to the local computer, and if port is [inaudible]. If it's not [inaudible], it will fail and return the result on the screen, on the java console screen.

As I said before, it's hard to reproduce the initial script, but you can make an assumption that the script is measuring asynchronous timeouts when there is a call to the certain TCP port. And of course it works with any IP address, including RFC 1918.

This is normal HTTP, how you get the feedback from the HTTP request.

The interesting part is the explanation. There was a call to the ISP representative asking what's going on. Why are scanning my computer? The answer was, well, we're trying to protect end users to find out if they have any malware on their systems so we can do proactive user query and tell them that they have a problem with their local computer or local network and that it's done to prevent leaking of credentials, the user credentials, and also to collect indicators of compromised devices.

And of course, the claim is to protect users but it's Internet basics, basically. If you have an ability to scan the port scan and you have some valuable – if you have some valuable results from the scan, you also can see if you have an SSH running or Telegram or TOR or any [DPM] network. Or Telegram, for example, in Russia, which is banned in Russia but still works.

The question is that – which, really, I don't know all the answers to this question. Is it really, the operation, in the user's interest? There are some thoughts about these findings.

First of all, TCP port probes on local host, not only local host but also the local network behind the firewall, basically. Even though most of the browsers have this cross-script prevention system but the

sophisticated way it's implemented and the idea of measuring the delays between get and respond gives us an assumption that how the network behind the firewall can be scanned. At least they can determine which internal hosts have TCP port in the [listen] state, of course.

The question is how do we know it is there? Is it common practice? We found three major services in Russia doing this thing. Also, after Googling the Internet I found that this kind of script, these kind of tools, are available from one of the well-known big brands providing marketing and advertising services on the Internet. Basically, they sell it. They sell this tool to the sites so they can – how they call it? So they will know the behavior of the users without disclosing their personal data or whatever they call it, anonymous scan without disclosing the data in the big aggregated formats. Like, if you have a website serving a million users, you have the script. And this is a part of marketing measurement tool, somehow. So, it's available for sale is what we've found. It's interesting.

The question is, of course, how do we stop it? Can we? How can it be stopped? I don't know. And of course this is a presentation without conclusions. This is what happens without the user's consent and likely without their knowledge. Most of the people have no idea what's going on behind the state. They use [inaudible] devices and web browsers and only for some reason some geeky guys find this, somehow found this interesting tool.

On the other hand, it's just a port scanning. Port is being scanned all the time This is also normal data [inaudible] ISP of any network. It happens constantly on the Internet and should we care about it? I don't know. Should we care if it happens inside our network? I don't know.

Also, we don't know if it's a common practice worldwide. All I can say that there are services for sale on the Internet. You can buy this thing and put it on your site and run this kind of script on the user machines.

So, it's a matter of trust, a matter of paranoia. I would like to hear your voice. What do you think about it? Is it a trust or paranoia?

Let's start with the trust. Who thinks it's a trust? Please raise your hand. Trust. You trust your bank or your ISP, so they protect you from the malware. That's what they claim. Okay. I don't see any hands. Who thinks that it's a bad thing and it's enough evidence to start paranoia? Yeah. See, we have a little bit special auditorium here. Most of the people here have a good technical background.

I was told by my colleagues in SSAC, they said, "Well, everybody will vote for the paranoia." I said, well, maybe we will find a single person who said we should trust our guys who scan us.

I like this little movie. I don't know how to start it. I don't know. I should click something, probably. That's a nice movie from Mars Attacks. "People, do not run. We're here to help you." It's a nice movie. [inaudible] friends. Please, we have people with questions, I believe, or statements.

EBERHARD LISSE: Thank you very much. This was frightening. But I'm quite sure we'll hear more about it because I really would like to figure out how I can check my banks, whether they're doing this, because I know their regulator. It's a small place in [inaudible], so if one of the banks would do that, it would be easy to handle. Anyway, Wes?

WES HARDAKER: Thank you. I'm Wes Hardaker from USC ISI. To answer one of your questions, you said, "Is this bad?" and the difference between port scanning on the Internet and port scanning this way is that it's happening behind firewalls and behind [inaudible] so that's I think highly [inaudible] on the bad aspect.

My question is did you do anything to try and analyze what's actually being done with the data, such as setting up an HTTPS proxy with a local cert or at least looking at the data, and if you turn on more ports, does that change the data outgoing back to the server? Because I suspect that that's what's going on.

ANDREI KOLESNIKOV: Not yet. This requires further research. This subject just came from Facebook from my friends. It is not like there is some team who is researching this. But all I can say that the banks and the ISP and also some companies who is making this kind of software [for sale], they are openly saying it's a good thing because we're trying to protect you. Really. That's what they say.

Because everybody knows malware started a long time ago with, for example, the DNS open resolvers, like the first instance of the huge botnets, the open resolvers from the DNS. And for example, the ISPs, they took some measures trying to find open DNS resolvers and close them and notify the network operators who has them, who operates them, etc.

So, maybe for the good thing they do scan and see some malware on the user machine. They call them and say, "Please fix your computer," or whatever. We don't know. That's the thing.

KATE PIERCE: Good morning. So, I have several years' experience as a penetration tester. I can tell you this is a technique I've seen before for several years. Is it bad? It depends who you're asking. But, yes, it basically comes down to someone is running requests from within your network. The biggest threats that I see around us are the interactions with certain firewalls which will open up an inbound port when outbound goes out in some cases, particularly dangerous as universal plug-and-play (UPNP). And cross [inaudible] request forgery to things such as web browsers. Or to a local server systems management. There are many machine management solutions that you could actually send request to.

But in this case, particularly with the things they're looking for, what I would be the most concerned with is unmasking your internal network address. For instance, if I can connect to your TOR proxy, I can send a request through TOR and a request out through your firewall.

ANDREI KOLESNIKOV: And see the IP address.

KATE PIERCE: And I can correlate your TOR circuit and your IP address. Also, if you have a high- to large-scale net gateway, I now know which specific outbound ports are associated with your sessions. This can get very bad.

ANDREI KOLESNIKOV: Okay, thank you. That needs further research definitely.

EBERHARD LISSE: Never mind that I'm much more worried about that they read my passwords.

ANDREI KOLESNIKOV: No. It's not about the passwords.

EBERHARD LISSE: Not just yet.

UNIDENTIFIED MALE: This is [Abdul] [inaudible] from Egypt. This is actually related to – java script is not detectable by vulnerability scanner. [inaudible] before publishing [inaudible]. This is my question.

EBERHARD LISSE: He asked is it detectable by a vulnerability scanner.

ANDREI KOLESNIOV: It's not detectable by the vulnerability scanner, you said?

UNIDENTIFIED MALE: No, is it detectable by vulnerability scanner?

ANDREI KOLESNIKOV: No. At least antivirus software is not giving any signal but we haven't tested all of them. It looks like a normal behavior of the browser even though it's a long story of the browser manufacturing to prevent the cross-scripting features in the browser, but still this kind of tells, still artifacts still there because it's a basic functionality of the browsers. They have to work properly. It's a game, cat and mouse.

EBERHARD LISSE: So, we got it hot off the press, further research required. Let us know when you want to present the follow-up.

ANDREI KOLESNIKOV: Okay, thank you.

EBERHARD LISSE: Okay. Thank you very much. And can you please stop frightening me? Ondrej Filip is the next in the same with DDoS Prevention at Speed. He is also going to frighten us.

ONDREJ FILIP: Of course not. I have two reasons to be nervous today, because first, this is the last presentation before lunch which is always a complicated position. Second, I am taking notes of this meeting and I'm sharing my file with Jacques Latour and I have no control what he is doing there and I see he is typing some things and it makes me nervous what's going to be there after my presentation.

Hello, my name is Ondrej Filip. I am from DotCZ, domain of Czech Republic or what we call now [inaudible]. It's a short name. We have never had that short name for [inaudible] countries, so we are trying to promote this. So, if I [inaudible] translated to Czech Republic, [it doesn't'] sound good for you.

I would like to talk about something we did in my country or [inaudible] in order to prevent DDoS attack. We took a very stupid approach. So, this presentation is going to be … Although it's about technology, it's going to be very simple. So, I think you will understand it easily.

So, what does it trigger for our work? [inaudible] DDoS attack against domain name of Turkey, DotTR. This attack was presented at Tech Day in ICANN 60. And there was one interesting sentence which I remember since that time and it's one ISP reported 220 gigabit attack bandwidth.

We are [inaudible] this on our side and he said, “Guys, that’s a lot. That could cause problems.” So, we have, of course, [inaudible] bit more farther and we decided that the current Anycast set up that we have is not sufficient, that we will probably not … We will probably have huge problems in [inaudible] kind of attacks.

So, we started to work farther. We also evaluated some hardware, too. Some traffic scrubbing centers and stuff like that. And those machines are amazing – amazingly expensive, also. But amazing if you scrub, for example, HTTPS traffic or all the TCP traffic. They are excellent. But for DNS, we always could find some way how to go through them because they [inaudible] is not perfect. DNS is not easy protocol and [those tools] were not helpful.

So, we are thinking, well, what can we do next? How can we improve this? Then, we decided what about just increasing the capacity of Anycast cloud? We started massive upgrade of [D] Anycast cloud. And of course we …. The DotCZ domain is most important for Czech so we started in Czech Republic and we did some upgrades there. And then we continued for neighboring countries and also to the last [inaudible] for Anycast.

Talking about Anycast [inaudible], the map, how Anycast clouds look like. So, the most concentration is inside Europe, of course. But as you can see we have some remote sites. Both coasts of North America and also South America. Also, we have one site in Japan.

Now let’s look more detailed into Europe. As I said, again, most of the sites are of course inside Czech Republic but also we have some sites in

Sweden, Stockholm, Austria. The new one in Milan. Then Germany, [inaudible], and London links Internet exchange.

So, basically, we are trying to put all servers to all the biggest exchanges because the best possibility how to pair with the address, I think Jacques talked about it quite well.

Also, when I say service in Czech Republic, [inaudible] there is certain exchange points which connects both countries, Czech Republic and Slovak Republic, too. So, basically, let's say we have all the connections – the best connections is Czechoslovakia.

This is the map of Prague. Not important, but just to understand that we have five sides in different [inaudible], connected to different ISPs. There is one side which is missing on this map because we have one hidden side which we don't disclose the location of it. There is not just DNS. There is also a [cert] running instance of the registration system and we don't disclose it because it's outside Prague and should handle all the traffic in case some big trouble inside Prague which we would not like to happen in this lovely city but we're already in for some massive problems inside the capital.

So, that's the situation. As I said, we started rebuilding all those nodes. We didn't add anything new except the [inaudible]. But we tried to upgrade all the nodes to a certain level and started in Prague and started to change those – especially those two nodes there into nodes that are connected by 100 gigs of traffic and with a lot of service.

So, the so-called big [stake] which we used in Prague has [inaudible] and it's connected 100 gig to the exchange point and has, let's say, also a very quick connection to the ISPs [inaudible] 140 gigs or something like that. So, in case of a huge DDoS, we have enough capacity inside the country and also we are able to serve the others.

So, this is the recap of the nodes. 17 locations, 10 countries in four continents or regions, if you wish.

And now how the process went up in the last two years. So, we started in a position that we have just 15 servers in Czech Republic and 16 abroad. We quadrupled the capacity of servers inside the Czech Republic. We went from 15 to 68 and more than doubled the capacity abroad. This is more an issue of cost, of course. It's much cheaper for us to increase capacity inside. We have [inaudible] there. We have good relationship with this, exchange point and so on. So, for us, this is much easier. Also, we have the technical staff there, for them rebuilding that side there is easier. Going abroad is more complicated. But also we work on that quite intensively.

Also, we started a new project, like putting service to some major ISPs in the country. They pay for hardware and for the service but they [inaudible] instance of DotCZ inside the network, So, again, even if we would be [dysfunctional], those ISPs would have a source of DNS traffic, at least one node. So, they would be operational as well.

So, this is how it looks in the number of servers and this is the improvement in the bandwidth capacity. We did some measurements. We are also DNS [inaudible]. We did [all] measurements. How much

traffic can we handle? It is true that we would almost be able to serve this capacity if we would get so many [inaudible]. So, we are very close to network bandwidth, by the way.

So, inside Czech Republic – and I can't count. It's almost six times more, right? [inaudible], again, is ten times more. So, we really upgraded all the connections to all exchange points and so on. Again, we did some connection in the local ISPs, the major ISPs.

So, this is the limits. This is just a recap. We now calculated that with the current setup, we could handle more than 200 million queries per second, which again might not be super effective in case a massive attack but we believe it's not so bad for this time. We can deliver more than 400 gigs of DNS traffic. We believe that's a good answer to some massive attacks.

Still, I know there are even larger attacks these days, but since this is a distributed system for the [inaudible], it would be pretty complicated to completely saturate all networks. So, that's the current limit.

We also thought not just about performance but also about diversity. We have everything multiplied. For example, one side – and I will show pictures later – is based on Cisco and Dell machines. Another is Juniper and HP. We have diversity in software, operating systems. And although we [inaudible] also DNS [inaudible] using the other software as well just to keep diversity because we think it's important. Every software has a bug and we, as developers, know it better than anyone else. That's why we keep two solutions for everything or trail. As I said, we are connected

to different networks, IXPs, and so on. So, trying to make this Anycast cloud as diverse as possible.

This is just a fancy picture just to impress you. It looks perfect if you go there. This is one of the nodes [inaudible] computers just doing the same stuff, basically not doing anything much because the traffic is ore than ten times slower than the node is designed for. 100 gigs [routers], so this is first – I think this was Juniper and Dell. And there's a second one [inaudible] white, blue, and red so it looks like a little bit of national [inaudible]. Again, this is based on HP and Cisco. So, those are the [cores] but we are building others outside the country as well, so that's just the beginning of the process.

Of course, now we have quite a lot of capacity, so we are sharing the infrastructure with others. [inaudible] is using that and Jacques can confirm that. [I know it's] in Macedonia, Tanzania, and Angola. Those countries are using this infrastructure, too. So, it would be probably illogical to just keep it for us, so we are sharing this capacity with others. So, if you are interested we can talk about it, of course.

Then, we build this and we were starting to think, okay, at least we believe you are kind of safe but is this cloud up to [inaudible]? That's another question, right? Who [inaudible] servers? What's their round trip type and so on? So, we did classical [inaudible] measurement.

Three years ago, we did ping measurement. We were pinging all the locations worldwide and trying to guess our TTN, stuff like that. Today, we use a different approach. We did passive analysis from the DNS data we have. We are building a system which is called ADAM, Advanced DNS

Analyzers and Monitoring software. Something which is not part of this presentation but I'm happy to share it next time, for example. And this is an output of the system.

So, we captured 14 days of traffic and it was something like 15 billion queries of [UDP]. And there was some TCP. And TCPs are important because we use them as a basis for our TT calculations which explains this graph.

You know how TCP works. At the beginning, there is a packet called [SIN] and [SINARC] and [AXEL] – three packets. And we can measure delays of those two when we send [SINARC] and [ARC] returns to us. So, this delay we called [RTD] in this presentation and in the TCP handshake and we collect this [RTT] for every client, every network. IP addresses. We store IP addresses [inaudible] and so on. So, it's quite a massive database which we collect from all those servers. Then we could make some calculations.

What we do, we calculated the median RTT for every IP address that queries us. So, this is an example of a client. So, this client approached five of our servers, three in Czech Republic, London [inaudible] and Vienna, and sent many queries. I can't [inaudible]. But many queries. Some of them were also based on TCPs. So, from TCP queries, we calculated the median RTT and now we know roughly what is RTT of this client.

What we do in the next step, we just do a weighted or evaluated RTT which is a weighted means of RTT for all servers based on the number of queries. So, we know that for this client the average RTT is roughly

18 milliseconds. Not so bad, right? So, this we calculate for every client that is connecting us and we have some TCP data. We don't have TCP [inaudible] but for many of them. This is the way how we calculate how well we serve this client.

Then we use geo-locations to figure out where those clients are. So, this graph – and I'm not sure how the colors work here – should show the amount of queries sent from a different country. As you can see, there's a small, very dark dot inside [Czech here] and the other countries have a little bit of lighter color. For example, US is quite important for us. There is a lot of queries. Germany, Russia, big country of course.

This is the way how we evaluate the distribution, where the queries are coming from. From this, you can judge not much except that we should probably work more in the countries that have the darker colors. But more importantly, we also need to calculate the RTT. So, another map shows what is the effectiveness of our DNS clouds where the RTT is slowest. It's quite what you would expect. Based in my country, countries around central Europe, northern Europe as well. When you are farther from those, the color gets darker. A

Again, this graph says the situation [inaudible] Africa is not [inaudible]. Also, east Asia. The RTT is quite high there which logical because you saw how the Anycloud is constructed.

But this is not a reason to put service into Africa because we need to combine those two information and I don't have a fancy map. I have just a graph that shows all the subregions [inaudible] and how well they are served. So, this is [inaudible]. So, to the right you are, the more

queries that region or subregion is sending to us and the [inaudible], the higher RTT is.

So, basically, those nodes that are on the right inside should be as slow as possible. And [inaudible] that we have some space for improvement, especially in eastern Asia. It's clearly visible. North America, [inaudible], southeastern Asia, and Australia. Eastern Asia especially are sending us a lot of queries and the RTT could be better. We need to improve there.

So, that's how we measure now the effectiveness of the cloud and how we play with that. This is the same for countries, this is top 50 countries. Again, you can see United States the RTT is not so bad, but China, Singapore, Taiwan, those countries are sending a lot of queries to us and RTT is quite high. So, that's probably a logical next step forward for Anycast cloud expansion. And you can see Europe, the blue dots, they are really low. So, they're pretty efficient. And there is a lot of people querying us as well.

One of the last graphs. I have thousands of those so I just put one as an example, that we evaluate every single country and then we evaluate which Anycast node is addressing that particular country. On the top, you can see Czech here and the biggest dots are in the nodes that are based in Czech Republic. So, this works really well.

There are some queries going to US nodes which is [inaudible] or to Brazil but they are really tiny, so not a big deal. IPv6 works as well. For example, Russia, the next line, the biggest dot is for Frankfurt. So, they are sending queries to Germany more than Czech Republic. There are reasons for that. So, nothing bad because it's still okay. You can claim

that Russia is big. So, saying Russia is a single country is a little bit complicated. We have even better breakage so we can tell you RTT [inaudible] in Moscow, for example. So, we measure those but I didn't want to be too detailed. So, that's how we organize our cloud and something we will work on. As you can see, they are not sending much traffic to [US].

And this is just a fancy slide. I don't know if it's good for something but I put it there. This is the representation of … The size of the country means how many queries it's sending to us, and the color, again, means RTT. So, as you can see, the biggest green spot, that's Czech Republic. The whole Europe is smaller than Czech Republic. It's something logical. The most queries are inside a country, but you can see US is quite big there and then you can see Russia, China a little bit. So, most of the queries goes from those sources.

And now let me conclude. So, we decided to use very silly algorithms, just to increase the capacity to a certain level. Now we are working with our software, too, called ADAM just to organize all the queries and make the Anycast cloud more efficient. It's a process we are still expanding. We are looking for partners in North America, Asia-Pacific, also Africa. It's something we need to work on. But there is not many good exchange points or it's quite complicated to [start there] but there are some exchange points growing so it's not so bad.

Just last remark. We build this traditional solution just based on PCs and traditional routers and now [inaudible] software. We are also doing some hardware development. We are now evaluating creating our own

DNS hardware, something like combination of [PGAs] and many [inaudible] based service. So, make Anycast like the big cloud, the big node, smaller but much more dense based on [inaudible]. So, if you are interested, it's something I'm happy to discuss. But again it's too big for this presentation.

So, that was all. I think we can go to lunch.

EBERHARD LISSE: Thank you very much. In a true Americanized fashion, throw money at it, especially when you are non-profit and you don't know what to do with the money when you are growing too fast. No, but we had this discussion. It's an ongoing—

ANDREI KOLESNIKOV: You know I disagree. Can you imagine what would happen if the national domain would break for the country?

EBERHARD LISSE: Yeah, but it's an ongoing running gag between him and me that we always say something about what to do when we are non-profit, lower the price or invest in human resources to develop this kind of thing. But my view is you can not keep up with criminals other than by just throwing brute force and just throwing money at it, invest into infrastructure. These people have lots of time figuring out how to counter things like this. This java script thing from banking is already

something that has made my day already and I must go and get my blood pressure tablets.

I like this thing about ADAM. You are already on the agenda for Cancun. Don't worry. I won't forget this time. Any questions from the audience?

Okay. So, then, I will release you for served catered lunch. We will return at half past 1:00. Mark Svancarek from Microsoft is in the room. He was conflicted out so we set him up for a time that he would be available, so he will be at half past 1:00 and talk about – sorry, not Mark Svancarek. Blanchette from [inaudible] is going to do the RDAP deployment. He is in the room, so we know he is there. We have received all presentations so all presenters are there. Enjoy your lunch. If you leave your things here, you do it at your own peril.

[END OF TRANSCRIPTION]