SAN JUAN – Tech Day Part 1
Monday, March 12, 2018 – 10:30 to 12:00 AST
ICANN61 | San Juan, Puerto Rico

UNIDENTIFIED MALE: Tech Day Part 1, 10:30 through 12:00. Monday, March 12, 2018. Room 209-BC.

EBERHARD LISSE: Okay. Good morning, everybody. My name is Eberhard Lisse, as you may or may not know. I'm the Chair of the Technical Working Group and the Managing Director of the ccTLD Manager for .na in Namibia. As usual, we do Tech Day on every ICANN meeting on Monday mornings. We have got quite nice agenda but before I even go into this, we want to see how many ccTLDs are represented and how many are gTLDs.

Can please – the ccTLDs from each ccTLD – one raise their hand, please? Eleven because Nigel is not raising his hand. And how many are the gTLDs? That's interesting. Okay. So, the rest are Technical Community, Security Stability, and whatever other SOs they fall under. You might be a subcontractor to the registries in the room running everything from DNS to backends or whatever, as well. Is there somebody who falls into this category?

The idea is that we want to see… We are still doing this under the umbrella of the ccNSO but feel it's more cross-constituency and we always wanted to move into supporting small gTLDs if they want that support, but maybe we must do more outreach. They probably don't know or they couldn't be bothered.

Anyway, that said, thank you so much. Let's quickly go through the agenda. Patrik Faltstrom has been briefed or instructed or asked or begged by Mr. Marby, who wrote a letter to the ccNSO about emojis in domain names. So, he's going to give us a briefing. We asked him to do that.

Wes Hardaker has got a setup at home where he probably has a local copy of a root server or something like this. Didn't really go through this into detail because I wanted to read the presentation and I was a bit preoccupied, so he's going to talk about that.

Then Mark Gaudet from CIRA, from .ca is talking about DNS as a defense layer.

If we could have the next slide, please. Lunch, unfortunately, is still self-catering. We haven't managed to weasel into the good graces of whoever sponsors DNSSEC Day.

Then we also have a host presentation or almost always. Pablo Rodriguez or one of his colleagues is going to talk about their

setup and, in particular, how it was affected by the hurricane. Given that my business partner died over the weekend, we are starting to look at business continuity. Our ccTLD was not affected because we had started to think about this getting older about a while back already, so we have processes in place, but this is maybe something that we need to look at on a more broader basis and a wider basis what do smaller or medium-sized gTLDs, ccTLDs have in place in case some catastrophe that is unforeseen happens. Are we prepared to deal with these emergencies?

At the same time, we will in Panama probably have a session about emerging threats. The SSAC is having some source into that, so we'll use this as a focus. So, if anybody has some interesting ideas about this for presentation and is coming to Panama, please note that we're going to ask for presentations.

Then Warren is going to speak about KSK Sentinel and then we will hear from the Loon Project. They're not mad. Those are balloons, in case you guys don't know. I saw on the TV two days ago about a drone being fired up by a cell service provider to provide a flying cell phone tower, so this is more for Internet connectivity but given the situation, we found ourselves in here or the locals found themselves here, it's quite pertinent and I'm very grateful to Google to make this [inaudible] available.

Then not on Hollander but Mark – what's his surname? He was sitting here just now. Mark is going to speak about universal acceptance and then we had a cancellation on short notice, so I'm very grateful for Jaromir Talir to offer a presentation about the open source registry software that they're using, which is called FRED.

Merike Kaeo is going to speak about IDN abuse and then Jacques Latour is going to speak, give us a follow-up about his Internet of Things security framework that we heard about last meeting.

The transitional break at 15:10 is with intent and malice because the SSAC has a conflicting session about – what was that again?

UNIDENTIFIED MALE: It's on name collisions.

EBERHARD LISSE: On name collisions and I didn't pay attention when the program group to which I observed discussed this, so I got basically stuck with having a conflicting session. That's not really a problem. I obviously had to try to keep the presenters, who will go to this meeting, away from presenting during the same time, and also, of course, I don't want anybody to leave but if anybody wants to

set up a little break that doesn't disrupt proceedings so much. Patrik?

PATRIK FALTSTROM: Yes. I also, as the co-chair of that working group, I also want to say that there will be an open SSAC Working Group meeting tomorrow. So, the cross-constituent meeting this afternoon will not be the only time when you can interact, so you don't have to go to the meeting if you find this being more interesting. The work party itself will meet – and I'm looking at you, Cathy – in 101B tomorrow from 8:30 to 12:30.

EBERHARD LISSE: In all fairness, I was prepared to actually locate the same meeting in our premises. I'm giving away the topics but then we had enough presenters that we arranged that we do the name collision. This is not ideal but I would pay more attention next time and fight harder that this doesn't happen again.

That said, sorry. One more thing. The thing is we have a remote audience. This thing is live streamed via the Adobe Connect. We are presenting from the Adobe Connect with one exception. The remote participants will see the same presentation as we do for we had a request from one presenter not to do this. I don't want

to go into this in too much detail. We're very grateful that they're coming, so we're going to argue too much about it.

After each presentation, there will be a little bit of time for discussion and remote questions will have preference.

PATRIK FALTSTROM: Thank you very much. Patrik Faltstrom, member of SSAC, no longer the Chair of SSAC. I will – is this one I'm using? Cool. So, I've been asked to within SSAC trying to coordinate the work that we do with internationalized domain names and related issues, where a few SSAC people, which work on these issues.

One of the things that we did last year was that we released document #95 with [a number] of recommendations. I will talk a little bit about what the document includes, but it should be sort of relatively well-known for the people that know anything about internationalized domain names, and then talk more in general what I believe personally the situation actually is somewhat, what's happening with it.

So, if we look at emojis in DNS, which is a question that's coming up, it is implemented just like all the domain names and all characters used via set of rules. And of course, the rules is that this sort of normal motherhood and apple pie that IDNA specified by the IETF and you cannot use emojis. But of course,

that argument doesn't really work, doesn't fly. People say, "But wait a second, if I use a domain name, it works anyways." For some definition it works.

So, but it's still important for us in the ICANN community to remember where those rules are set. They are set by their IDNA standard that is developed in the IETF and by the IETF. And if we look at details regarding emojis and some other symbols is that they are of the Unicode category symbol other, which also not IETF defines is actually Unicode Consortium that defines what code point is in within each category.

So, Unicode Consortium has defined that emojis are symbol. IETF have said symbol others are not good to have in identifiers. And ICANN has said we follow what the IETF is saying, so the three organizations are sort of following each other.

But anyways, the people just say, "Oh, but we can ignore that do things anyway," but that is sort of the formal situation. SSAC did take a look at this and even though you can just read the standard and spec and say this is actually not possible, the question that is, "But what is really the problem?" And there are a couple of issues that I will bring up here.

The first one is that emojis are very visually similar. They can be very difficult to distinguish them. The difference between some of the emojis are much more different between different

operating systems than the difference between each other. And this is because one could see the differences between the emojis in, for example, Android and iPhones as being font variations for other characters. So, it is very difficult to know what emojis is which one. One example is that there is Unicode Consortium. If you look at their tables, there are more than 20 different emojis, which are sort of smiley face or variations thereof.

The other thing that SSAC was looking into, which is not really fun, has to do with composition. Because there are certain emojis that can be glued together with a Zero Width Joiner, so for example, a man together with a woman together with a boy ends up being a family. The question that is, of course, like there are no rules that are set on what's happening if you do a combination of other things. For example, what happens if you add a second boy, would that a family of four or will it be a large boy, a very tall one like Liman. Or what happens if you have… yeah.

So, with the composition, you can see like I'm a mathematician at the bottom. You see you can create all different kind of number of permutations and it's a little bit unclear what's actually happening if you do these kind of compositions. That's not really fun, not when you talk about identifiers. We have to remember we're talking about identifiers here, not drawing pictures.

And then we have this modification because in Unicode 8, they introduced the ability to apply a modifier than say that it will actually apply different skin colors to different emojis. For example, by adding one of the skin tones to the code points. But the question, then, of course, is what really happens if I change the skin tone of horse or a house or a car. Is that how you change the emoji to be a blue car instead of a red car? And what if you want more than five skin tones? So if you had two skin tones or if you do that to the family? Or… yeah. Anyways. Undefined territory. And when talking about identifiers, me myself, I don't really like the word undefined.

The next thing, which we in SSAC believe is quite serious, has to do with accessibility. Because the emojis are visual constructs and there's no standardized way of speak or enter an emoji via voice, and this ends up being an accessibility issue for disabled people. We think that is something that should be taken very seriously, specifically when talking about identifiers.

So, in the report, if you go to the report itself, this is what the five findings that we have in SSAC, they are disallowed by the IDN standard, they are not required by designed standard convention to be visually uniform or visually distinguishable from each other. Modifiers and glue arrangements of them via the Zero Width Joiner, for example, allows for potentially a much larger set of composed multi codepoint symbols.

We also found that when you have two domain names, which are identical in appearance except for ordinary typographic style variations, but have different underlying codepoints, which is what ends up happening here, they actually will by definition identify two different DNS domains. And so far with domain names, when we talk about font and typographic style and stuff, we said that that is fine to have a domain name in a red font or green typeface or red typeface. That doesn't change the identifier but with emojis, it does. That's a very big difference.

And finally, from an SSAC perspective, we do believe that we need to continue the IETF inclusion policy for including codepoints that we allow in the identifiers we use in the DNS. So because of that, it's just straight unrealistic to expect that just because a codepoint is including Unicode, it should be used as part of a domain name.

That led us to the recommendations that we have two of them. Because of the risk, we recommend, we repeat an early recommendation that ICANN should sort of, of course, but we don't really say that in writing, of course, not accept TLDs, which have codepoints that are not valid in IDNA 2008, which means we should not even start talking about using emojis in TLDs. And the second recommendation is that we strongly discourage registration of any domain name that includes emojis in any of its labels because of the issues that I just described and more.

And we advise registrants of domain names with emojis that the domain names that they might have already registered or tried to use or whatever, that they may not function consistently or may not be universally accessible as we expect with domain names.

After we released this report, there are a couple of things that has happened and mentioned a few of them but specifically, there was a Board meeting on 2nd of November, 2017, and there were four different result statements there. Here are two of them that ICANN is now acting on.

The first is that the Board directs that conformance to IDNA 2008 and its successor will continue to be a necessary condition to determine valid IDN TLD labels. And from SSAC, this is something we have been waiting for since, I think, 2002 or something. So, it only took about 15 years to actually get this crystal clear from the Board and I think that's a good thing.

The second resolution is that the Board requests the ccNSO and GNSO engage with SSAC to more fully understand the risks and consequences of using a domain name that includes emoji and of its labels and inform the respective communities about these risks. And one initiative that started is that ccNSO when we and SSAC have engaged with each other and we will have a discussion on Wednesday in the ccNSO workshop, where we'll

actually present the same slides, so for those of you going to that one, it'll be kind of boring, this is part of the outreach that we are doing.

Regarding the first one of these results, let me personally comment on that one. What this one says is that ICANN Board is actually recognizing the work that is done in the IETF, and I think this is really, really important. And this is a cooperation that is recognized between the standards organizations that many of us have been looking for and I am personally very, very happy to see this engagement because it says very clearly that ICANN Board does not see ICANN being the organization that determines what is safe and what kind of characters and codepoints should be used in the DNS protocol that is something that IETF decides upon. So, for those that would like to have something else than what is recognized, please go to the IETF. That's one way of interpreting the first one of these results.

Some more personal comments – last slide – that I think people should be aware of, which explains a little bit how massive the situation is. On the other hand, when I made this slide, the situation was very messy and I was sort of maybe in a bad mood day and had not had enough coffee. A lot of things actually happens at the moment, so all of these things are on the way to be resolved in one way or another.

The first point is something that I would like to make people understand and that is that the Unicode Consortium runs and adopt-a-character program. It literally gives the ability for people to pay Unicode and get whatever character they want. Let that sink in. It's a funding model for Unicode Consortium. You go to them, you suggest what character you want, you pay them money, and you get it. Yes, exactly. Is that similar to a process in some other organization we know something about? By the way, it's much cheaper to get a character than a domain that in TLD.

Well, there are many. It gives sort of some interesting flavor to how characters get added to the available characters that we're using. To me, it definitely says that no, we probably should not allow every character and codepoint that is added to Unicode. I think it's a good thing that we are saying yes or no and have an inclusive program there.

Anyways, we have early versions of Unicode 11 is released for comments. Please have a look at those. Please apply the IDNA 2008 rules, which are the normative ones, to that and see what your findings are. Let me know. I am the [inaudible] appointed expert that also have a look at how the IDNA at the result of applying IDNA 2008 rules on the Unicode versions, what actually happens.

Unicode 11 includes a bunch of new emojis, including a lobster, and you might have seen the tweets from a governor in one of the U.S. states, which are really, really happy that we finally have a lobster in the Unicode codepoints, because then we can have that as a domain name. He could have got one just by going to Unicode and pay for it. Anyways. So, that includes maybe that will create some pressure, political pressure of using emojis.

It's also the case that there are various parties that including global domains international, they had a presale of domain names that include these emojis, and that's a kind of interesting business model and specifically as so the characters do not exist in Unicode yet and on top of that, they are not valid according to IDNA 2008, and we still see a market. We see sellers and we see buyers of these domain names, which I think is an interesting market economy situation. I think we in the community discussed this situation.

161 single emoji of Unicode 11, these codepoints that do not exist anymore in a version of Unicode that do not exist was actually sold. On top of that, they were, as we was talking inside, SSAC, we said emojis are so confusable, so not even the registry could keep them apart because they are actually sold the same emoji into multiple domain name, multiple registrants, which was interesting.

That kind of glitch, of course, is something, a problem, a mistake that anyone can make. I'm a programmer myself. Oh, I should not talk about what bugs I have in my software. But anyways, it's kind of an interesting situation that when a party tried to sell domain names that did not exist, that are illegal, they also failed. But it happens. We're all engineers. We know that can happen.

But what is even more interesting, if we look at W3C, if you look at skin tone and those kind of modifications, at the moment, they are looking at adding the ability to have those modification as part of a style sheet. And let that sink in a little bit. Remember that different skin tones on emojis create not different fonts. It creates different characters in Unicode. Okay? So, to be able to know what character you're going to have in an identifier, if it is the case that whatever the text is in the document is changed due to the CSS, you need to have a CSS parser to actually determine what codepoints we're talking about. That attack vector regarding confusability I think is really interesting and could make me sleep really bad at night.

It is created for it to be text decoration but they are talking about also using those kind of modifiers for arbitrary characters to be [overlaid] on others, and that implies, of course, the ability to add modifiers to codepoints, which is the skin tone, for example, or create a family, let's say, "Oh, all Unicode emojis,

please add a boy to all of them," or "I want skin tone white or yellow on all the emojis here, including the poop." I don't know what that means.

Anyways, that's all. Thanks. Questions?

EBERHARD LISSE: Thank you very much. Give me a good hand. Any questions? Please, identify yourself for the remote audience and for the transcript.

[ABDALMONEM GALILA]: [Abdalmonem], ICANN coach. I have one question and [inaudible] by comment. Let me think out of the books. Could we standardize the shapes across different operating systems and across all devices to have same Unicode and same shape? This is the first part of the question. What is the challenge to do that? Second part. Is this will make DNS to accept emoji to be [existing online]? Third part of the question.

My comment. Using emoji will bring more people online. I think that. Thank you.

PATRIK FALTSTROM: Okay. Let me try to understand the questions and give an answer to that. First of all, you're talking about the… I think, you

ask about the actual design of emojis, whether we could harmonize that across devices and displays. That might be possible to do. I don't really know where that work should be done and if it is in W3C or somewhere else. I see it is as somewhat problematic just because… Sorry, let me take a step back.

Yes, it would probably be easy to have it more harmonized. We've seen the discussion regarding the cheese, whether that is above or below the meat on the hamburger emoji, and it might be possible to harmonize that, but it's not done in here in ICANN.

It's also the case that I think you are talking also asking question of whether it's not the case that it's possible to select subset of the emojis and say these are safe and not include all the emojis, specifically maybe not the ones that everyone is paying for and get their favorite crayfish or something to emojis.

That would imply in the standard that IETF start to add individual codepoints to be allowed in the Unicode, in the IDNA standard, and that would practically be work that is done in the IETF. The IETF has so far said that they don't want to do a character-by-character evaluation what is a safe character, not safe character. It is complicated enough with, for example, the Arabic script that is used in languages that are non-Arabic, and those kind of problems are things that need to be resolved.

Regarding your last comment that more emojis could bring more people online, I might agree with you on that but I also think that it will also disadvantage everyone that cannot read and type, for example, people with disabilities, so I think that's a big danger. I also think that we should focus on getting all the languages in the world available in domain names, for example, all the languages on the southern part of India and in Africa that uses the Arabic script. I think that is far more important and that will bring many, many more people online than what would adding even a single emoji will do. Thank you.

EBERHARD LISSE: One more question.

ELAINE PRUIS: Hi. I'm Elaine Pruis. I provide some consulting services to registry operators and registrars, some of which are considering offering emoji domains. So, you may have already answered this but I'm wondering if the SSAC or IETF has a safe list or a danger list, and if you're working on that or maybe you just said that you don't want to undertake that.

PATRIK FALTSTROM: The answer is that that does not exist and emoji is just like other symbols or not including the IDN standard and we do in the

review that we did in SSAC, we see enormous number of risks to use emojis. And once again, one of the problems is that a typographic change. It has to do with the design of emojis compared to other characters in Unicode code set. A typographic change of the emoji changes the codepoint. It is similar to as if you change from one font to another one, it will be a different identifier, and that will be a bad thing if a domain name is different if you use it in Time or in a serif or a non-serif font. So, just the whole design of emojis need to be completely different to be something that can be used.

So, to me, it's a complete nonstarter. That said, if it is [inaudible] that work should be done going down this path because people want to use various symbols as part of identifiers, that work must be start in accommodation of W3C, Unicode Consortium, and IETF. We don't have the competency in ICANN to do that more than, as you say, and as we hear, and as we see, potential market economy forces that would like to do things. But there are many things we would like to do.

ELAINE PRUIS: Okay, thank you. So, you also said that –

EBERHARD LISSE: Sorry. We are running over time, so I don't want to have another question. Take it offline to talk to him or raise it tomorrow on the next meeting, please. I don't really like to shut down discussions but we're running a little bit ahead of time. It's much more complicated than it used to be 20 years ago when I created my own little emoji with a stethoscope in ASCII.

Anyway, Wes Hardaker is the next presenter.

WES HARDAKER: Thank you very much. So, I'm going to talk today about a project that I have at ISI called LocalRoot. It's sort of in beta at the moment and it's definitely looking for more – oh, clicker. Yes. That would be good. Thank you. And I'm looking for people to play with it and tinker with it and see if it's helpful to you, so we'll go over what it is.

So, it is basically a project that lets you run a recursive resolver with precached root data, and eventually, it'll have other data, too, for other zones, and we'll get more into that in a minute. So, classic DNS resolution, which I think most people here understand, well, the important part is the caching part of it. So, as you are, say, sending out www.example.com from your laptop, it goes to your ISP's resolver typically, and then it goes to all portions of the infrastructure tree.

And note that the green common example.com end up getting cached so that the next time you want to go see a different domain. Maybe you go to ICANN.org. Well, none of those are in the cache it ends up adding org and ICANN.org to the cache.

But if you end up going to exam.com, for example, the com portion is reused. You don't go back to the root in order to go figure out where com is again. So, this is standard DNS hierarchy with caching.

So, what I'm trying to do with LocalRoot is for a zone that you might want to have precached because you use it all the time, and the root's an example of something that's used all the time. RFC 7706 from the IETF talks about doing caching of the root when you're in bad networks or in bad places. This kind of takes this farther and says let's just always cache it.

Within ICANN, there is sort of a push to do. There's another project I think within ICANN called the hyper LocalRoot, where they're trying to do other ways to distribute the root data, then over DNS directly, and so this is just an extension of that.

So in the end, if you are using my LocalRoot project, you would have the root zone precached in your resolver ahead of time always, and what that does is it means you basically never contact the root. You never need to anymore, so if the root went

offline, which has never happened. Let's say it did. It actually wouldn't affect you.

The one big difference that LocalRoot provides over getting data in some way, because you can actually just transfer the root zone and do this yourself. You don't need it. You can transfer it over http at a [cron] daily or out of some timer. There's a number of ways that you could do it and then to substantiate the file yourself into your resolver. This gives you notifications, so the instant the root's changed, the notification will arrive to your resolver and your resolver will be a request transfer for like AXFR or IXFR, one of the other mechanisms for doing it.

As I said, ICANN actually has a project to try and start pulling root data from other mechanisms, as well, like over http directly within a resolver.

So, what's the benefits of doing this? A couple of things. One, you get sort of the pseudo-caching of the root data. Caching is really not the quite the right term there. You're actually running an authoritative copy of the root data for your local network, but we'll just call it for caching for ease of understanding. It removes the need to contact the root. I've mentioned that. You get faster DNS lookups for the TLDs. Because of caching, that's not going to provide a huge amount of benefit. There's some suggestion that NX domains actually help quite a bit, especially if you're on

a low-latency network, to know that when somebody starts typing names and your browser starts looking up stuff early, that you'll actually get answers faster that show that that domain that they're typing doesn't exist.

And then always, you have an always up-to-date copy of the root in case you get connectivity issues with whatever with root system, you won't actually need to contact them.

And most importantly, I'm trying to enable research of your own. So, you can trigger events after a root zone notification. Since we will be sending you DNS notifies, you can actually do stuff with that metric to go do analysis or whatever else. I'm actually looking for people that are interested in coming up new ways to make use of the notification system, as well.

So, a couple of things about security. You do know that the root zone is signed, DNSSEC signed, so that's actually what enables projects like this or any other mechanism for fetching the root data because it's signed. As long as you check the signatures, you can get it from anywhere that you want.

LocalRoot, the way that I've designed the interface, and you'll see pictures of that in a minute, does do data transfers with a TSIG that's generated for you so that the connection between the local root server and your resolver is actually secured with TSIG. Because it's DNSSEC-signed, that actually doesn't provide

you. It's like a double benefit. It's like a double-check. It actually doesn't provide you a whole lot of extra security over DNSSEC itself, if you validate the data on the resolver side.

So, I'm not going to do a demo. I'm going to do screenshots instead because it's easier with Adobe Connect. So, this is what the website looks like and it is localroot.isi.edu. You're welcome to go create accounts and sign up yourself and I'll repeat that again. It's localroot.isi.edu. And you go through the process of typically doing the normal thing. You register with an e-mail address and a password and state that you're not a robot, all those types of things. You'll get the e-mail note saying that you can log in now once you click on a link, and then once you're in the system, there's some links.

I know these are small but the next ones will be bigger. The important one is the getting started link actually walks you through what you need to do, and there's only a few steps. Basically, you create a TSIG key and there's a link for doing that. You create a server and then it will spit you out some config, as we'll see in a minute.

So, this is the TSIG generation screen. You click create new TSIG and so TSIG is, for those who don't know, is a mechanism for protecting DNS traffic using a shared key. And once you create a TSIG key, you can name it whatever you want. It will show you

the key and it'll show you the value. You don't need a copy of the value or anything. You just go on to start creating a server, and when you create a server, you're going to give it a name, you're going to give it an IP address, and then it's going to ask you what TSIG key you want to use.

In the future, some people have suggested "Do I have to use TSIG right now?" It's all locked together. And yes, you do. But the reality is some people may not want TSIG enabled, and that's fine. I will make that optional in the future.

And then you have a server list. So, again, there is my server. It's running at 10.0.0.2. Obviously, not an address that works on the Internet. It's using my cool TSIG key. It says it's enabled. It still says the active status still has an X in it, and the reason for that is that I don't actually select it as active and it actually send you notifies until we see an initial transfer from you. So, actually, when you create the config and actually stand at the resolver, one of the first things it'll do is try transfers, and so you'll actually get a transfer quickly or you can actually run dig to LocalRoot to actually get that from the correct IP address to get active immediately. And I do that to make sure that you're actually using an address that you own so that when I get the transfer request from you, I'm not going to send notifies to addresses that you don't own.

Note that in the little green line on the far right, it actually says, "Get config," and so one of the things that you do, if you click on that, you go to this configuration generator. And right now, it only does bind but there's a couple of different bind options and so this is what it looks like if you select full recursive resolver configuration, meaning it gives you an entire recursive resolver config set that you can just then go throw into bind and you're up and running, that's it. And in that configuration set, I'm not going to scroll down it, but it includes all of the servers that you can transfer root data from. It includes your TSIG key that we just generated, and it binds the TSIG key to the local root server, and everything that you need, basically.

So, it's trying to be as easy as possible for somebody that doesn't necessarily have a whole lot of experience setting up resolvers. At this point, because of its beta status, I'd prefer users that know what they're doing rather than somebody that doesn't have a whole lot of experience running a resolver.

So, what happens when you run it? Well, this is a graph of real world effects and that big downturn right in the middle was where LocalRoot was turned on, so this is a recursive resolver. There's not a huge number of requests. This is running in my house but I'm a geek and I get a lot of mail and I've been running a mail server for years, so I'm actually sending somewhere between 5 to 40 requests per… I think this is in a minute window

to the root and actually… yeah, I think it is in a minute. Anyway, to the root, and so I actually send quite a few, especially because of incoming spam and doing reverse DNS lookups and stuff like that. And that gap in the middle was where LocalRoot was turned on and you can see it fell too close to nothing.

The little tiny blips in the bottom are actually not real requests. That's because I also have a RIPE Atlas node because I'm a good RIPE Atlas user that likes to contribute data to the Internet, and those RIPE Atlas nodes still send queries directly to the roots. So, this is actually measured between just me and B root, by the way. This is not to all the roots.

So, future work. As I said, this is sort of beta in status. It's been up actually for three or four months now. I've added a bunch of stuff recently. It's actually being monitored much more than it was in the original, so I'll be notified when there is trouble. I have other people that have offered me VM instances so I can stand up other servers. Right now, there's just a single DNS server that you can slave from.

I'm going to add other zones, so rootservers.net and ARPA will come very quickly, actually. Probably maybe even by the end of this week so that you can mirror those, as well. And then I'm looking for partners that want to offer other sources of zones

that people really care about or that you might want to do, as well.

I have a request for students at ISI to help deliver data from that you can get from other TLDs, for example. Obviously, TLDs that are huge, you don't want to do this with .com, not that I can get .com's data, but you would never want to run that in your local resolver. That's way too much data. But somewhere, there's a balance of there's possibly some ccTLDs that are actually quite small in general that would be worth pulling in, as well.

I'm looking at other transfer mechanisms, especially with ICANN is looking at that, so I'm going to let them do it and hook off of them when they get there. Some variable period update notifications. Because the root actually doesn't change that much, the interesting thing is that you don't really need to pull the data that frequently. You need to pull it on a regular basis for DNSSEC purposes, but you don't need to pull it twice a day, if there wasn't any change. So, I'm considering playing with sending notifications only when there's change or when on a regular basis, that's less than twice a day to get you to update your DNS signatures because they're going to expire at some point. So, some fraction of the root zone signing period, and that will be that configurable variable on your end.

And then the other thing that I'm thinking about adding is some monitoring error reporting, because it's sometimes hard to tell if this is running. When it's in your resolver, unless you're monitoring your traffic, it's actually hard to tell whether your resolver is self-querying itself for root data. So, I'm actually thinking about doing some monitoring at least from some of the sources to say, "You know what, you're still sending data to the root." We're seeing it through other mechanisms and that way, I can be sure that so I can actually tell you yes or no, whether your service is working or not, or let you know when that status changes.

Any questions? Please let me know what you think. There's the address again. I'd love feedback for things that you might think might be useful additions or if you have research focus that you think that this might help with or you need features, hit me up.

[JOHN LEVINE]: I'm [John Levine] from [inaudible]. Hi, Wes. This is a dandy idea but I'm confused. How does this differ from RFC 7706?

WES HARDAKER: It's very, very similar. RFC 7706, if you read the fine text in it, says you are not supposed to do 7706 unless you have a valid reason for doing so. So, basically, if you're running on a lossy network or

a poor network and let's just say that the 7706 author told me explicitly I should never call the 7706, so I'm not doing that.

[JOHN LEVINE]: Okay, 7706, even though I do DNSSEC validation, so I don't worry too much about TSIG and stuff like that and it looks fine. Observed fact is the root zone changes twice a day, even if it doesn't change, they update the serial number.

WES HARDAKER: Right, no. I said that, but that change is actually useless to you to a large extent because there's no data changes, so the change… If I look into and I assume you're referring to the variable period updates – there's a couple of reasons for doing that. One is that if you're on a really bad link, you actually don't need the new data because the DNSSEC signatures are valid for a long time and the new SOA record won't help you too much.

You need some changes on your side, too. That's a longer-term project. The other thing is that you may want notifications only when data actually changes from a research point of view rather than going to look yourself. I may actually be able to send DNS notifies to people to say, "Hey, by the way, something just changed in the root that wasn't signature or SOA-related and wasn't boring, in other words."

[JOHN LEVINE]: I presume my client is doing [inaudible] so the actual amount of data is low. I understand if it's a lossy network, like a reconnection is painful but like in the normal world, how much of this is sort of is a practical advantage and how much is just it's cool?

WES HARDAKER: That's a good question. Some of it comes from… Were you and OARC or no?

[JOHN LEVINE]: No.

WES HARDAKER: So, I've given two presentations in the last month that talks about privacy aspects and one of the ways to make sure that you never leak data is don't ask questions. So, if you have all the answers ahead of time, there's actually a privacy benefit there, it turns out.

[JOHN LEVINE]: Even though I have a small network, I think there are – I ignore Warren's advice and I slave the root anyway.

WES HARDAKER: There you go. So, if you want the notification, you can sign up for here. That's up to you.

JEFF OSBORN: Jeff Osborn, F Root. Hi, Wes.

WES HARDAKER: Hi, Jeff.

JEFF OSBORN: I got to ask this. If this is wildly successful, do you think it would have any impact on the global root server system and do you feel like commenting on it?

WES HARDAKER: That's a good question. I mean, so for those who don't know, I'm actually the operationally in charge of USC's root service. So, in some extent, I'm actually taking myself out of the business if this becomes successful. I doubt it because it would have to get deployed by default in CentOS or something like that, and I would have to work with vendors that would want to do that.

But there's a lot of work going on, so one of the rationales for doing this, there's a lot of thought of are there alternate

methods for distributing root data besides the way we're doing it now, so this is a toy to experiment with that. And if it became wildly successful, we'll all need to work together to make it wildly successful.

JEFF OSBORN: Yeah. You might look into planning for crazy success. It's happened before on the Internet.

WES HARDAKER: It has, and I have thoughts around it but… yeah.

PAUL: Paul [inaudible], CentOS vendor, I guess.

WES HARDAKER: I said that because you were in the line.

PAUL: So, how is this different from AXFR [inaudible] still?

WES HARDAKER: So, really the notification is the only major difference. So, you, as I said before, you can actually functionally slave from the roots today. Because there's a list of roots that you can do AXFRs from

today and, in fact, when you get the local root config dump, it actually has all those, too, so you can't just pull from us. We actually give you a whole bunch of other places so that you can do this safely. If any of them went offline, if LocalRoot died, the config would still work because you're still getting data because you'd still do SOA checks in order to figure out whether your data was up to date and that kind of stuff.

So, with respect to serve stale, serve stale will work up until some combined period of time. There is no difference.

PAUL: But if all the root servers would just open AXFR, we have like what, 6,000 points we can AXFR from? Like I'm not sure how this is really different.

WES HARDAKER: It's not significantly different. This makes it easy to do. So, there's only a couple of things. I said this in the beginning, right? That people already do this. You can already do this with your zone. That was like one of the first statements I made. You can dump the root zone config and slave it from any of the roots that allow AFXR, which there's a few, plus ICANN servers that let you do it, so that's easy. You're done. This provides notifications, which you don't get through the other mechanisms. Your server

has to do SOA checks and things like that, but who cares? That's like an hour delay or something like that.

This makes it insanely easy. Most people don't know how to add that zone. I actually dump config for you so that you can just copy and paste it in. It's an ease-of-use thing, too. Some of it's publicity, right? I want people to start slaving. I want people to start doing AXFRs and this is a way to make it easier to get more people to try it and experiment with it.

DUANE WESSELS: Hi. Duane Wessels from Verisign. So you had a slide that talked about benefits and I get skeptical about projects that only list their benefits. If this is all benefits, then why aren't we doing it? So, are there any downsides to doing something like this?

WES HARDAKER: Yeah, there is. I mean, just let's take… I lost the benefits slide. So, that's a good point and I should create a list of downsides. The most obvious one being if you are not a resolver expert and you tinker with config and you don't succeed in a transfer, you don't realize that you need to deal with large TCP transactions through a firewall, you're running a new service, a new copies of something, so you better know what you're doing because you

could break your resolver. And so if you're going to do this on a major ISP, be careful.

You could always do it in one resolver in your ISP and not with another, if you want to tinker with it for a while, but no, that's a valid point. I'll work on creating… There it is. I'll work on creating a detriments slide. Risks would be a better word probably.

[ABDALMONEM GALILA]: [Abdalmonem], ICANN coach. I assume that I am the resolver at [ISP] so I will exhaust myself to looking for the cache for the query. And at the same time, I have the local root zone at my side. So, I think there is no need for cache now or somehow as a client, maybe have the same local root zone and no need for cache. By the way, I don't trust [ISP] so it's meanwhile for me to have local root zone at my own machine.

WES HARDAKER: Yeah, you can. There are mechanisms for standing up local resolvers on your infrastructure yourself, be like laptop or server or whatever. It's more difficult on some operating systems than others, but in many times for things like in a Linux-based desktop, you can just install [name] DNN and it's up and running pretty much and copy config from here, so it's very easy on some

instances and not on others. I don't know how to do it on Windows because I never have but I know it's possible. Warren's going to tell me how to do it on Windows.

WARREN KUMARI: Oh God, no. So, Warren Kumari, Google. I was one of the original authors of 7706 and a few people have said that I suggested against doing it unless you had a really good reason. I don't think that's quite correct. The document kind of says that because the document had to say that to get published.

WES HARDAKER: Right, and so I wasn't blaming it on you. I wasn't blaming it on the authors. I was saying that the authors were told that had to go in there.

WARREN KUMARI: And so yeah, I mean, this is basically 7706 with notifies and automated magic, so this is I think wonderful.

WES HARDAKER: Thank you. Now I'm in trouble.

UNIDENTIFIED MALE: Very quick comment. You mentioned TSIG not much benefit. I think one thing is that it gives you slightly stronger guarantee of freshness, man in the middle.

WES HARDAKER: Yeah. I mean, there's other things. That's true. It does help you in a couple of other ways, especially because most resolvers, most authoritative engines don't check the DNSSEC validation, so unless you turn that on, a lot of them just take its own data from their master and serve it anyway, so [inaudible].

DAVID LAWRENCE: David Lawrence. I just wanted to say that I don't have a very good reason for doing it and yet I'm doing it. [inaudible] public declaration that if you want to stick the IETF RFC police on me, I'll be your test case.

WES HARDAKER: Well, thank you for doing it but appreciate it.

UNIDENTIFIED MALE: This is [Simone], Salesforce. So, I just wanted to follow up on the TSIG point and this is something I pointed on the RSSAC Caucus mailings recently. TSIG does provide additional protection for transfers of a full zone because remember, DNSSEC by itself only

signs authoritative data in the zone. The root zone is full of nonauthoritative data, such as child NS delegations and glue records, so if you want to be sure that you have obtained the root zone correctly, you need a channel protection mechanism and TSIG is the sort of DNS protocol resident channel protection mechanism, so I would recommend that if you're going to do 7706, you need TSIG or you need some alternate way of verifying it, such as https. You mentioned that as a possibility, too.

WES HARDAKER: Yeah. It's a very valid point. Thanks very much.

UNIDENTIFIED MALE: We have an online question. If someone wanted to learn more about ICANN's hyper LocalRoot project that you mentioned, I believe the answer is talk to the ICANN Office of the CTO, but if you could specify.

WES HARDAKER: I don't know the right contact when I believe that would be a good place to ask because they would know where to point you. I don't think that that project is at the point of deployment yet, so I think you'll hear more about that once their side code projects are done and completed and deployed. Good question, though.

EBERHARD LISSE: Okay. Thank you very much. Give him a big hand. One of the rare cases where short presentation generated more debate than the length of the presentation, which is very good. And because I like a little bit of interaction that we don't get fall asleep all the time.

Next one will be Mark Gaudet and CIRA and talk about DNS as a defense layer. And we are five minutes ahead of time, so no problem.

MARK GAUDET: Okay. Thank you very much. My name is Mark Gaudet. I'm a Product Manager at CIRA and I'm responsible there for DNS services that we launch. We run some TLD DNS as well as an Anycast DNS service within Canada. And what I'm going to do today is talk about, share our experience launching a DNS security product in Canada. So, it's a cloud-based DNS firewall. Share with you what we learned launching the service and what our vision is for the service going forward.

So, cybersecurity, like everyone in the world, is a big challenge in Canada. We recently did a survey of roughly 2,000 businesses in Canada, and from the feedback and results from 1,000 businesses, the results aren't surprising. Most organizations use firewalls, they use antivirus software, but what is surprising is

19% of the respondents from the survey said they were victims of a ransomware attack, so one in five. As well, 32% of the respondents also said they were a victim of a phishing attack where they actually gave up information from inside their network to hackers.

So, there's a major gap. There's still room for improving security, and one way to do that is through DNS. So, DNS is used by every application, malicious or good. 91% of malware uses the DNS for command and control. So, by looking at the DNS, you can have a view of what's happening in the network, what malware is operating by observing the DNS queries. The good thing, it covers all devices, doesn't require any new software. It works for everything. It gives you a view of all traffic within the network.

So, one of the ways to leverage DNS to build a security product – this isn't new, there are products out there – is with a DNS firewall. So, at the core of it, in the purple is a policy-enabled DNS server. So, on the right is the corporate network sending out or customer network sending out DNS queries. The policy-enabled DNS server examines all the queries and compares them to a threat list, as well, look at the responses, the IPs that come back in the DNS query response to see if those are on the threat list, as well. So, prospective a mechanism there at the DNS level to examine and block DNSSEC queries. That's at the

enforcement, the ability to report what is blocked and why, providing information that could be used to mitigate.

We have approximately 250 organizations using an authoritative anycast service, so they started to come to us asking about recursive DNS and if we could offer a service. So we in that exploration, we started to, we realized there was a requirement in Canada, a need for a secure recursive DNS service that could also do content filtering.

We looked at building that ourselves. We looked at buying, we looked at doing some custom development, but what we quickly realized was that the real value in a DNS firewall or any other security product is the threat fee. How is that threat fee produced? How dynamic is it? What is the data used to produce that threat fee?

So, we realized that we weren't going to be able to. We looked at public feeds, commercial feeds we can buy, and what we realized was whatever we produced was not going to have much incremental value from services that were already being used for security, so we needed something that had a different threat fee.

We went out into the market and looked at different companies and what we ended up partnering with Nominum, which has now been acquired by Akamai, so I'll use those interchangeably. So, Nominum sells recursive DNS software to large service

providers globally, and what that gives them access to is large amounts of DNS data. So, they answer in the orders of trillions of DNS queries. They take a subset of those queries and get an anonymized feed from all of their ISP customers and analyze that to produce a threat feed.

So, that was the real reason we chose Nominum was that they had access to data and the ability to produce a high-quality dynamic threat feed, and they add an average of 100,000 domains to the threat feed each day. So, this just I could spend a whole presentation on just how the threat feeds produce but in general at the top, they have the service provider data, which is roughly a million DNS queries per second coming into their data science center. They combine their data science with commercial and public threat feeds to produce a dynamic threat feed that we receive at our policy-enabled DNS server, which is Nominum software.

So, one of the things they focus on is newly seen domains. So, if a domain has not been queried in the last 60 days, it's automatically suspicious and is analyzed, and within 14 minutes of it being seen for the first time, if it is malicious, it will be added to the block list and we will start blocking it. So, real value is that it's really dynamic.

Just to give you a little bit of information on the service we launch, we've taken the Nominum Akamai software and built a cloud service in Canada operated by CIRA. We launched it in June. We had 80 organizations live on the service at this point. Behind those 80 customers are more than 500,000 network users. That's because much of our base is in the public sector, large universities, school boards with tens of thousands of users in each of those networks.

We're currently running at peak of 3,000 queries per second, so starting to generate a lot of traffic. On an average month, last month, what we're growing, so in February, we blocked a million malicious domains. The threat feed is very dynamic and over that month, 5 million new domains were added to the threat feed.

We do trial, as well, and 100% of all networks that have turned on this service and have seen blocks, so there are things in their network that have been identified as a threat from the threat feed. As well, we've talked to lots of organizations and they still see a gap in their existing security solutions, and leveraging DNS as an incremental layer of protection has been very well accepted in the organizations I've talked to.

One of the things that was a concern was false positives. So, that hasn't been an issue. For the number of users we have behind

this service, we might get an average of one or two queries per week, which are "Hey, this domain is blocked. Why is it blocked?" And each time we do an investigation, there is a valid reason for it being blocked and a lot of confusion happens when a customer has a site that they've commonly gone to and used. An example would be an education, a school board where someone was downloading educational materials one day, the next day the site's blocked, and it was because that site had been compromised and was now distributing malware. So, it requires some level of support to analyze the sites that are blocked.

I'll shift a little bit to what we learned in launching the service and operating that since June of this year. Some of the operational requirements, customer experience, and benefit, as well what we receive as a view on the threat landscape in Canada by having access to this DNS data. And as the registry for Canada, what benefits has brought us in being involved in this service and operating it.

So, one is operational, a recursive DNS service is open to attack. Our service is whitelisted for only the organizations and networks using it, but we still see attack traffic. This is a graph of blocks related to mostly pseudo-random subdomain queries that have been blocked by the service. These are relatively small given the capacity we have, but it is a fact that it will be attacked

and we have functionality within the service from Nominum Akamai that does detect pseudo-random subdomains and blocks those and applies rate limiting to different malicious traffic directed at the service.

As well as a cc operator, we have no experience in operating a recursive service. We know what kind of traffic we'd expect at the authoritative level but not at a recursive level. So, this graph shows you from the time of launch the average query rate over 24 hours, so starting out in the hundreds to over the past month, we're averaging 1,500 queries per second and with a peak of 3,000, roughly.

What we've done to try to figure out how many networks we could support and what types of networks we can support is to do some capacity modeling. So, this shows the number of network, how many queries per second, per thousand network users. So, at the top of school board is generates a fairly low number of queries, so that's like one to two queries per second per 1,000 children at a school. So, we can support lots of school boards.

At the bottom, we have some small ISPs using the service, up to 10,000 subscribers. And the difference there is they don't forward queries. They send the queries directly to us, so they're roughly 30 queries per second per 1,000 users. And then you get

a range in the between cities using it, municipalities, universities, or commercial businesses.

So, what we've been able to do there is take our capacity of 200,000 queries per second and be able to map that into how many of each type of network user we can apply to the service.

From a customer perspective, they will see a dashboard that shows them their DNS query traffic plus the number of blocks for each type, so in the middle, the orange is malware, is phishing domains in particular. Domains that a user has clicked on and then on the right is botnet blocked. Those would be malware acting and sending out command and control DNS queries, for example.

And then this gives the counts by threat type. If you look on the right, the threat names, broad categories of threats, which include devices being used for bitcoin mining, ransomware, different types of malware that are found on a hospital network over the period of 30 days.

And one of the values of this service is it provides a view of the threat landscape in Canada. We're just starting out but with 500,000 network users, that's a fairly spread geographically across Canada, that gives us an interesting view on the threat landscape in Canada, what's happening, what threats are appearing.

So, this is an aggregate view at the bottom, the different colors or the different types of threats that are shown here from January to the end of February. And if you… showing in the next slides, I'll turn off. So, one of the things we quickly saw and when we started looking at the data in January was every single network that was using our service was, we were blocking bitcoin mining, so sites that were using user PC resources to do bitcoin mining. That really peaked in January, and then it dropped off. So, lots of we could dig into that data further and try to figure why it dropped off. Was it starting to be blocked by other security products? But it's very interesting. We'll see things happening very quickly. We saw that before it was in the press and visible.

This shows just blocks associated with Mirai. There's a steady state. If I extend this into March, you'll see a steady state of queries. A lot of these are more in an ISP and home users, but then there's a big peak at the end of February.

General, some of these are grouped into more general categories, so malware call home. See there was a something popped up at the beginning of February, where there was a large number of queries.

So, we really get a good view of all of the malicious queries that are being blocked, which provides information on what's

happening in Canada, and that provides lots of interesting opportunities to do some research on the data, as well as identify new threats specific to Canada, so we have access to the data or starting to do some data science on it, where we would start to look for things that aren't on the threat feed that are specific to Canada.

It provides us an opportunity to inform and educate. So, when we see specific threats that are occurring, we can provide outreach to Canadians in general, Canadian organizations in general, or to users of our service to tell them about specific threats that are showing up, what they can do to prevent them. If it hasn't hit them, here's what is happening across the country.

We also, in discussions with organizations using it, universities, for example, that use this service. One of their biggest threats is phishing and it could be we block phishing domains but some could be very specific to Canadian universities. There's an interest among the university community to share information, so if they identify specific phishing domain, we'll create the ability for them to share that information to quickly block it on our service and to have that protection from that more quickly.

And it also has really changed CIRA's role in Canada with respect to cybersecurity. We're starting to be in discussions with other groups in Canada within the government and in not-for-profits

that produce threat feeds where we will incorporate their threat feeds into the service to add more of a Canadian flavor to it.

So, going forward, the vision for… we started out this as a primarily commercial service. There's some we are trying to diversify revenues, so we do have an interest in selling it, but we're also seeing that now that we have access to the data and we are operating this service, there's a lot more we can do with it and we will push it to share that information, keep data sovereign, as well, in Canada. There's an interest in keeping the DNS query data in Canada and answering those queries there.

As CIRA, a not-for-profit, we aren't going to commercially sell that data. We will keep it private and there is a trust level there for us to use it to improve security in Canada. As I said earlier, we'll augment that with specific Canadian threat data that we can get access to, and look to provide longer-term leverage this technology to provide more of an open service for home users all across Canada.

So, some of the things just to conclude. The partnership we came up with Nominum and Akamai really helped us move forward to quickly. So, we've taken what is an enterprise carrier grade service and deployed it very quickly within Canada and built a cloud service there.

DNS as a layer of protection, as one layer in a defense in depth strategy is really gaining momentum. There's more than 50 organizations that are paying for this service as a commercial service, so they believe that it can incrementally improve on their firewall and existing security.

And we see there's a benefit as a CC to operate this service within Canada to own the DNS data query data, and to be able to leverage that for research and to improve the security of Canada's Internet. I see the DNS data as a huge valuable resource that we can leverage. And that's it. I'm happy to take questions.

EBERHARD LISSE: Thank you very much. I think he deserves also an applause. And I abuse now the prerogative of the Chair and ask question from the [inaudible]. A small ccTLD like .na, have you basically… When we see this, we manually blocked slash 21 sometimes just to punish the owner. How can we use such a system?

MARK GAUDET: We can take that offline. So there is –

EBERHARD LISSE: No. Don't take it offline. I'm not the only manager of a small ccTLD and I think it's a pertinent question. We don't have the financial resources to contract with big providers. Fortunately, Nominum is one of our anycasting secondaries, so maybe we can talk about it. But I would like to hear what you suggest for small cc or as a gTLD, especially ccTLDs in developing countries.

MARK GAUDET: One of the ways we could help out as CIRA is that this is implemented as a cloud service. It's easy to extend it into another country, so we could host a recursive service that's fed by the threat feed in another country and make that available and make it part of our operation. That's one of the reasons I'm here presenting this is, A, we'd share our experience and, as well, if someone was interested in deploying this, we have the capability to help them do that.

UNIDENTIFIED MALE: Okay. So, this service is I think clearly aimed primarily at being the last hop iterative resolver, right? So the applications go directly to your resolver and you see what traffic they're asking for. If I put another resolver between the application and your resolver, things like query minimization come in and you might not always block and look up at the zone apex, but if you want to go deeper into the zone, you might not see those queries, they

might go to an authoritative server, if there's another resolver downstream from you.

Question. Do you also envision serving intermediate resolvers downstream, RPZ, redistribution, or anything like that?

MARK GAUDET: The way the service is typically deployed now as most organizations, all of the organizations using it have a recursive server inside their network and they've configured it to forward the queries to us, and they would lock down the DNS to only allow queries to the service, so they would cache locally and forward to us.

UNIDENTIFIED MALE: Okay. So, you do the minimization for them but they expose the full queries to you.

MARK GAUDET: Yes.

WES HARDAKER: Wes Hardaker, USCIS. Unless I missed it, you didn't define what actually blocking does. I assume you're returning an NX domain to –

MARK GAUDET: If it's malware, like a nonuser, if it gets back a SERVFAIL, if it's a user, like a phishing domain, for example, there is a proxy server as part of this service, so it will get redirected to a blocked page to let the person know. It also does content filtering, broad categories of content filtering, and that will have a content focus block page.

WES HARDAKER: Yeah. I was thinking specifically in the context of a DNSSEC-enabled validator sitting behind that was forwarding through your resolver and what that would look like to them, but a SERVFAIL would still end up being a SERVFAIL in any case.

UNIDENTIFIED MALE: [inaudible] with DENIC. As you obviously are in the position to identify and to block domains, don't you fear that you are being asked to take them down and seen as a censor of a net?

MARK GAUDET: Not at this point. So, all the organizations using this, they own their network, they're turning it on and choosing to block and forcing acceptable use within their networks. So, we have not hit, and as well, we're not doing active research yet on the data

we receive to identify anything new. We're receiving the block list from Akamai and enforcing that, so the censorship is already being done in the networks that we work with. They're already using firewalls to block domains. They're doing content filtering, they're doing custom blockings of specific domains, and this is another mechanism they can use.

EBERHARD LISSE: Any more questions? I would like to repeat the exercise we did earlier when we started when the room was only half full. How many ccTLDs are represented here now? Don't tell me it's less than before. How many gTLDs are in the room? How many? Two. How many subcontractors for gTLDs and… [inaudible] together with the gTLD please. The obvious question is we want this to cross-constituency, so is there something… Please think about it. If there is something that you feel we can do to engage more or to increase the percentage of gTLDs, especially smaller new gTLDs.

Maybe take this out with you for lunch and you can contact us at ccNSO Tech Daily e-mailing address is open so that you can basically post to that list. It's not a closed list. We have a separate address for internal business but if you posted that list, you are allowed to post and [inaudible].

That said. Thank you very much. Go and have lunch on your own expense.

[END OF TRANSCRIPTION]