BUENOS AIRES – Tech Day
Monday, June 22, 2015 – 10:30 to 17:00
ICANN – Buenos Aires, Argentina

TOILEM GODWIN: I received a call from one of our registrars, who is a friend of mine. He told me he could not access some of his .ke domains. The first thing that I did is I checked everything on our end. I didn't know it was a DNSSEC issue, so I just confirmed DNS is working okay, everything is fine. Then I asked them for access to their servers, I confirmed the name server and everything. I started troubleshooting on their end, instead of confirming everything is okay on my end.

We normally do that because I find most registrars might not have that good knowledge of DNSSEC, and they don't have DNSSEC understanding of troubleshooting. Either of them are college students who want to make a few bucks, so they just buy hosting services. They don't even understand mostly, technically, how DNSSEC works. We usually assist them in troubleshooting some of these issues.

Now, since it was a major problem and I spent a lot of time troubleshooting an issue that was not there, I found the response team was a little bit smaller. I found instead of looking at the problem on the registry I ended up spending a lot of time troubleshooting on the registrar end. We normally do trainings for registrars, and because of this issue we've planned to have more registrar training so that at least when they have problems, we don't spend time on their servers, but on our servers.

After confirming everything on their registrar servers was okay, the problem was okay. I didn't understand why, so I started troubleshooting all the domains within my network. Within my network I have a few domains that are already signed, but I don't know, for some reason I didn't believe it was DNSSEC. That was part of the problem. We experienced a problem where for some registrars, some domains were accessible, and some were not. We normally check second-level domains, so we confirm with government, NGOs, those second-levels that they're okay.

We found all domains within our network first of all were accessible. I did not understand why then, but we came to later. Government sites were accessible. Although they are not on our network, those websites and domains were accessible. Later on we discovered that domains that were accessible had name servers that refreshed after more than five hours. Some were still accessible even after eight hours, after an outage or something. I think that again brings a question where if using DNSSEC, why were these sites accessible?

They should not be accessible if you had bogus records. Maybe it's part of my understanding that's wrong, but maybe some of the experts here can clarify that. We continued with the hard work of doing troubleshooting. Then after a while I came to understand that [Leminow 00:04:27] really do thorough checks within the DNS of all .ke domains, of the root .ke domains. Something told me, "Just check DNS, because maybe the issue is there. Maybe the reason why we have a problem is there."

But to me, the reason I didn't consider checking DNSSEC is because I had started the process of rolling over my keys. I knew the process of the rollover should happen in the next 30 days. I had already started that process, so I did not consider it that much, because the process had already started. But again, to my surprise, when I went and checked the keys, the signatures had expired. We did not expect the signatures to expire that day - we expected the signatures to expire within the next 30 days.

For us, the signatures expired earlier, and we've been trying to investigate why the signatures expired earlier, because we had already set the expiry of the signatures for April, not March. But I think we've done a bit of research, and we used a different server. When you are testing, before we deployed DNSSEC, we used different servers with different versions of BIND. We don't know which BIND had this bug, because the signatures expired earlier than we expected.

One of the things that we discovered was on BIND, there is DNSSEC auto maintain, which I think we set to "on" and with auto maintain the understanding is when a signature reaches an expiry date it revoked and even it deleted some of the signatures that we had. That was another mistake that we discovered on our name servers. So we continued doing troubleshooting. We wanted to resolve the problem. One of the big mistakes we did was we contacted IANA to remove the DS records, because for us, that was the easiest process, but it ended up actually being the longest because of the procedures that IANA has in updating records to the root.

When we started that process, again, because most of our email addresses were in .ke domains, and we had not updated the emergency contact like adding a gmail domain or .com address, we could not get confirmation from IANA. That again started the problem. We called IANA and they told us, "We've send you emails. You should have received a confirmation link." So that one did not work. We decided then, "Let's us find the keys," which is one of the things we should have done first. We should have found the keys that had those DS records and signed the zones, which was an easier process than contacting IANA to make those changes.

We normally do weekly backups for the keys, even if it's the same keys. We just do the backups and restore them somewhere. For some reason, we looked for keys for around seven weeks, and all the keys were corrupt. So when we sign a zone it doesn't work. Again, we don't know why the keys were corrupt, but we found one key pair that we could use to sign the zone. Since the whole process, I think we ended up being in a panic mode. You'll find that we signed the zones, it was okay, but now the problem is again we've skipped the procedures - the procedures you need to follow when signing the zones.

You are supposed to change maybe any records, like SLA records on a signed zone before you actually sign them. For us, that panicking mode and everything, we signed the zone and we changed the SLA records of the signed zone, and we got a bogus signature on the SLA

records. Anyway, when you get that one error on a specific record, we noted again you could be able to still access your .ke domain.

In my understanding, if you have any bogus records in your domain, I don't know if you should not access, but since when you're accessing a domain you're accessing through the NS records, maybe you can still access it because you're using the NS records and it's confirming the signatures? That's maybe another clarification that we may need from the house. Basically, we went through the right procedures and processes and we brought the signed records back.

Although after around 24 hours IANA had removed the DS records. For a while now, we've not updated any DS records on root. We have been doing a lot of research to understand what we really need to do. I think for us, we rushed into implementing DNSSEC before really understanding how you should be able to handle issues that come when you've already deployed DNSSEC. But to encourage other registries, I think it's a good thing. Maybe it's just you need to understand, and maybe if you have a team, let them understand how you need to resolve problems with DNSSEC.

After that incident, we created our own root. We normally do tests from assuming IANA on one server and everything, so we do the process, testing with a very rigorous test from root to the second-level domains. We are still doing that, and we have come up with scenarios. We can make a key invalid and then try resigning the zones, try to do a lot of tests with it. For us, the part that we still are looking at is that

part of the signatures expiring a month before the date that we put there.

I remember very well we put the dates on April, and even we noted it down that this was the expiry date. I don't know why the signatures expired a month earlier. Those are some of the things we're trying to look at. The versions of BIND that we've been using, somebody's looking at it so we can know if it's actually a bug that anyone else using BIND has experienced the same problem. We just had simple DNSSEC monitoring, and now on the tests environments we are setting up a detailed DNSSEC monitoring.

We are coming up with our own tools for monitoring DNSSEC that will give us details of that. The other things that we are doing to prevent issues that we have, I think it's best that you choose if you want to maintain your DNSSEC. We did not - and I think it's an honest mistake - we did not know if you put DNSSEC automation on it will bring you issues. I think it's good, if you put DNSSEC automation on, you decide if you're going to do it on manual or you want BIND to do the automation process for you.

Lessons learned - some of the lessons we've learned is the first thing we shall do, if DNSSEC fails in your network, is try actually recovering. Try and find your old keys, resigning the new zones, because it is a very short process. For us, we went with the second option, which would be the third option, where you are going to remove the DS records. That was one. The other was for us, because we have such a small technical team, I think we were so much engaged in the problem that we forgot

communicating with the other involved parties, so that is another part - communication is very important to the parties involved in those failures that you are having.

The other one we learned is that before you fully deploy DNSSEC on the environment, I think you should do a practice of all possible DNSSEC failures that may come. You should have a separate test environment where new failures keep on coming. Practice on those new failures and check the response time that you have, in case this failure comes, so that when you're communicating, you'll know the approximate time you'll be able to resolve a problem in.

Again, check the pros and cons of setting automation on, if you're using BIND. I've not used any other software but BIND, so I don't know if the other ones have this option. I think that is it for me today. Thank you.

EBERHARD LISSE: First question, from the Chair's prerogative, what software do you use exactly to do this?

TOILEM GODWIN: We don't use any software right now, for the .ke zone. But for our second-level we are using Open DNSSEC to do the automation of resigning the zones, but we are still testing it. We were actually supposed to deploy DNSSEC on our second-levels in April, but this problem happened in March so we had to push the deployment further back.

EBERHARD LISSE: No, what I meant is the software that failed, that caused this problem? What was that?

TOILEM GODWIN: It was BIND.

EBERHARD LISSE: Yes, but how do you run this? For example, in .na I do DNSSEC key sign offline, and then load the signed zone up to the primary. It's a separate process. It happens on a different server in our primary.

TOILEM GODWIN: For us, our set up was a little different then. We had a different machine where we signed the records, and we just uploaded to the active server. But now the way we've designed our network, we have now [unclear 00:18:44] servers behind the master, which is now the one that's accessible behind the other one.

EBERHARD LISSE: Right. Can all speakers identify themselves for the record? Because it's taped.

MARK ELKINS: Hi, Mark Elkins from DNS South Africa, Domain Name Solutions. A couple of questions. Are any of your keNIC domains DNSSEC signed?

Are you signing any of your own stuff like NIC.ke or something like that? If so, do you run a DNSSEC aware resolver for in-house use? Obviously you have a resolver sitting in the office somewhere. Is it DNSSEC aware? Wouldn't that be a good idea, and that you might pick this up a bit quicker because your own stuff wouldn't just suddenly stop working? I personally also install the DNSSEC validator on my laptop, the stuff that the czNIC guys did, which is absolutely fantastic.

My other thoughts are what sort of training processes have you gone through? We use Open DNSSEC because we think it's a bit cuter, but before doing that we had a look at a lot of the different stuff. You are aware we do training down in South Africa.

TOILEM GODWIN: Thank you. We have keNIC, .ke, that is not accessible outside. But we use it to test the DNSSEC on keNIC .ke. We have DNSSEC aware resolvers in our network. The last question, I did not get it.

MARK ELKIN: Yes, but surely if you had DNSSEC aware resolvers, and you had the little things like the DNSSEC validator in your browser, wouldn't that have started showing that there was a problem, and that it was the signatures that were broken immediately? I secretly went around Kosa House and got our resolvers to be DNSSEC aware. I then went around to individual people and said, "Wouldn't you like to install the DNSSEC validator tools?" - quietly, without making a big fuss, so that when

things break it jumps into people's faces quickly. My last question was looking at Open DNSSEC and further training.

TOILEM GODWIN: The validator on my machine, I didn't have it installed then, but after the incident I set it up on my machine. We do trainings on DNSSEC, and also we have been looking at other options other than Open DNSSEC. That is why we pushed deployment on second-level domains further forward. Because we had actually made up our mind to go with Open DNSSEC, but we wanted to understand how other options worked first of all, before we deploy DNSSEC on the second-levels.

JOE ABLEY: Hi, it's Joe Abley here from Dyn. First of all, I wanted to thank you for coming to describe the problems that you had in front of the room. I think everybody who has a problem should come and describe it so that I think everybody can learn, and I think it's fantastic what you did. The actual question I had was a simple one. Years ago I worked at ICANN and I was there when we signed the root.

We had this process where a TLD could say when they'd contacted the IANA, "This is an emergency DS record change," and if they said the phrase, "Emergency phrase," then suddenly there was a process by which the entire change could be completed in 48 hours, I think, or maybe less. But you have to say it's an emergency. I think the whole time I was there, nobody ever said that it was an emergency, so it was never treated as an emergency change.

I was interested if you told them it was an emergency, and the second thing is that for everybody else, just be aware that there is a process for emergency changes of DS records.

TOILEN GODWIN: Thank you. Yes, we contacted IANA with the request that it was an emergency. But I think when you look at it, if IANA is doing it within 48 hours it's better for us to resolvce the problem within ten or even two hours. That is why I said I think the changing of the DS record should be the last option. It's better you try to find your old keys and resign the zone to restore your records with the DS records that's on the IANA database.

JOE ABLEY: Yes, I agree that's the best way.

DAN YORK: Dan York with the Internet Society. I too wanted to say thank you for coming here. I hope more people will do this kind of thing. One question I had for you was basically, now you're in a testing process, you're looking at this - what are your plans to bring the domain back to be fully signed?

TOILEM GODWIN: We are planning in August to bring the domain back. There is still one pending issue we are looking at, and we are expecting that by then we should have sorted it out. Again, because we had to redesign our DNS

network, initially we did not have slaves behind the master that update the accessible slave, which other people see as a master. We had to redesign the network after the incident to give us at least other options of adding additional slaves, to make it more robust. We're still working on that, before we start the updating of the DS records to IANA, because that's the only thing that is remaining.

DAN YORK: Great. Best wishes with that. We've got a community here who's open to helping, so please do feel free to reach out to folks here if you need more help. Thank you for coming here.

EBERHARD LISSE: I'll take two more questions. Rick Lamb, and then after the second one I'll close.

RICK LAMB: Hi, it's Rick Lamb from ICANN Tech, or whatever everyone's been saying. Thank you very much for sharing this with us. If every one of these incidents from the .uk ones, .fr, all these things in the past, you're in good company, and it helps us all when you explain this stuff. Also to reiterate, yes, you are in good company. If you want any help, Mark Elkins here has said as well - plenty of training, plenty of opportunities.

As a last point, from my experience as a geek and following this stuff around, I love the BIND auto sign, I love DNSSEC, I love the developers, but I think what you're doing is good; to take your time to try and

understand how this stuff works. For a lot of cases I've seen the script works much better, because when you do have an emergency like this - and I'm glad you had that first point "don't panic" because you want to be able to calmly look at the thing, understand how it all works, and be able to get each piece out.

That becomes a little harder with something like Open DNSSEC that can solve almost any problem. That's just my personal view that I've seen. But thank you again.

TOILEM GODWIN: You're welcome.

EBERHARD LISSE: I'm taking one more question.

VICKY [REISS]: This is a comment more than a question. I'm Vicky [Reiss] from ISC. We're the BIND publishers. It makes me very sad to hear a story like this. A couple of things. One thing I would encourage is anyone who has an emergency like this to contact ISC at support@isc.org, and if you're in a situation like this where your TLD is down, we will absolutely help you, and we have all the expertise on how to make DNSSEC work on BIND. We have quite a few ccTLDs that are using BIND. We have a fair amount of experience with it.

That said, I know it's very complicated to keep it working, roll your keys correctly, et cetera. We wrote a New DNSSEC Guide recently, but I

know that both in our guide and generally, our troubleshooting information and recommendations is generally quite poor. I would really welcome you to take a look at that and give us feedback, especially with the experience that you've had, so that we can improve that for others.

EBERHARD LISSE: Okay. I can't take any more questions. Drew Bagley is the next presenter. Thank you very much.

DREW BAGLEY: Good morning. My name is Drew Bagley, and I'm with the Secure Domain Foundation. I work with Norm Richie, who many of you know, and for those of you unaware of the Secure Domain Foundation, we are an international NGO dedicated to proactive anti-abuse efforts through collaboration, information sharing, and the use of an API. Because our mission focuses on proactive anti-abuse, we decided to do research into the factors that affect business costs for different models of anti-abuse for Internet infrastructure providers. So today I'm going to present findings from our first research report as an organization.

We called the report "The Cost of Doing Nothing" because we literally wanted to see whether or not it was more advantageous to be reactive and do nothing about anti-abuse until a complaint comes through a registrar, registry or a hosting company, or if there perhaps were business incentives to be proactive. Setting out to do this research, we

were asking two primary questions, and crafted our research accordingly to find out exactly what the business costs were, with the whole picture associated with anti-abuse efforts, and then to specifically look at different models of anti-abuse to find out the pros and cons of each, from a business perspective.

To better understand this topic, we really wanted participation from all Internet infrastructure providers, but for this first report we focused on registrars, because that's where we got the highest participation. In drafting a survey and asking questions about different anti-abuse practices, we looked at the context in which registrars operate in terms of the legal responsibilities, the reputational incentives, and the financial incentives to act in certain ways with regard to anti-abuse.

In our survey we purposefully left our questions very open-ended so that in this first research report we could really get a broad sense of what type of companies we're doing, and really not make any assumptions as to what they were doing, and let them tell their story. Now, as is well known, there are obviously legal incentives to do something about anti-abuse, at least once a complaint shows up at a registrar. For gTLDs there's the fear of ICANN accreditation, to make sure there's something being done about complaints, and for ccTLDs there are different local laws, and of course community best practices.

Tying into this, we found that of course there are reputational incentives to come across as a clean registrar so that your TLD is not blocked, or so that your sub domains are not flagged for suspiciousness. Also, we definitely found that there were financial

pressures with regard to potential credit card chargebacks, for when you have to take down a domain name. Then of course there are lawsuits, court orders, and what-not that all add up and can create business costs. So it's within this context that we drafted our questions.

Of those registrars that responded for this first research report, in total they represented about 12 per cent of the domain name market, and were geographically diverse. We're hoping down the road we can make it even more geographically diverse and increase the representation, but for this first report this is at least a pretty good indication of some of the trends. We also had a range and size, going as small as 800,000 to up to 15 million domain names that were managed by the registrars that responded.

We found a lot of diversity in the methods used, which is exactly what we wanted. We wanted to hear from registrars that were merely reactive - only doing something once they received an abuse complaint - and also wanted to contrast that with those that were proactive and took action from the onset. So we were successful in at least getting a pretty nice, diverse range of respondents.

Our key findings really showed us that bad customers are of course bad for business. The more bad customers that you let through the floodgates that become actual registrants, the more it can add up in terms of business costs. We found that registrars that are proactive, and therefore are able to filter out bad customers from the onset, are

able to reduce the number of abuse complaints they get on the tail end with their actual registrants.

The most expensive part of anti-abuse efforts, universally, were the labor costs associated with responding to abuse complaints. With those labor costs we found that of course you can reduce those by reducing the number of bad customers, but also something that our respondents really told us is that they spend a lot of time dealing with complaints that are poorly worded, or perhaps even directed to the wrong parties. So better complaints can also be better for business by reducing labor costs.

I've temporarily lost my slides. Going back to the labor costs associated with anti-abuse efforts, we really found a range with which respondents said that it cost them in terms of hours, as well as in terms of money. For some respondents, they claimed to spend as little as 15 minutes per abuse complaint, but that still was 15 minutes that was fixed because labor costs can't really scale. Once you get to the point that you are receiving an abuse complaint, and you have to have someone from your legal team or whatever, or maybe even an outsourced abuse desk look into it, there is a certain manual review process that everyone who responded to our survey followed.

Therefore, that part could not really be automated and could not be scaled. The costs pegged to the amount of time spent too ranged greatly. We had the time going from 15 minutes, to some registrars claiming they spent an average of ten hours per abuse complaint, and then with costs it ranged from .80c up to some saying that they spent

up to $60 per abuse complaint to resolve them. Even with the registrar that responded claiming they had the lowest cost in dealing with abuse complaints, added up with the number of abuse complaints they were still spending at least $100,000 a year, and upwards of $250,000 a year to address abuse complaints.

What this showed us is that no matter what, even if you get down a pretty good system, once you get those abuse complaints it can still be pretty costly for a company. As I mentioned with the diversity of approaches that different registrars demonstrated to us that they take with anti-abuse, we really had the full spectrum. We had some who were so purely reactive that they literally did nothing, except for very basic financial validation, meaning that they made sure that the credit card would go through when a would-be registrant was creating an account.

We had others that would at least look into their resellers and flag things for suspicious credit card activity - once again, just tied to the financial transaction side of things. Then we started seeing some who were a little bit more proactive. There were some who would at least use specific key words in their filtering, so if a registrant tried to create a new domain name that had the word “PayPal” in it, then that was treated as being suspicious. So there were some steps taken to try and avoid the creation of phishing domains.

Then we had some very proactive respondents that from the onset use either our service, the Secure Domain Foundation’s API, or a similar service, to actually run a query of the registrant’s provided

information, to see if it's been associated with any known malicious activity. For those, we deemed those to be the most proactive, because they were actually looking at associations from the onset and truly trying to screen out bad customers before they became bad for them.

Each respondent described a very similar process when it came to how a registrants registered a domain name, to how an abuse complaint was filed and investigated. Essentially it's in step three on here where you see those activities really diverge from reactive to proactive. Once you get to the WHOIS verification stage, there were some who were very proactive, there were some who just made sure the email address worked, but did nothing else beyond that.

Then once a would-be registrant became an actual registrant and actual customer, then it was a wait and see game to see if abuse complaints rolled in, or there was ongoing proactive screening from some respondents who would continue to screen their domain names to see if their domain names, or the email addresses associated with their domain names had become associated with any maliciousness.

If you look at steps six, seven and eight, that's where you get a lot of those manual review costs I was describing earlier, where perhaps for some respondents they said they could get it down to as low as 15 minutes per complaint, or .80c per complaint, and for others it was on average ten hours per complaint. That's where on the tail end you have a lot of labor costs that can add up.

This is the part I was able to do off the top of my head, so you have this information. This visual representation I think really shows what our results told us. When we were looking at the reactive registrars compared to the proactive registrars, this is essentially why we started seeing a divergence in terms of the number of complaints they were getting on the tail end and the amount of business costs associated with each method, and the amount of money that could actually be saved.

As you can see, for those who engaged in proactive screening, that's where they were able to literally filter out bad customers that otherwise would have just gone through and been treated like any other customer, and become a financial problem for them down the road, as well as an abuse liability. Then those who did ongoing screening similarly could continue to week out these bad customers before abuse complaints came flooding through. This diagram really visualizes that and shows that.

Something that really stood out for us as a prime example contrasting reactive anti-abuse to proactive anti-abuse dealt with two registrars who should have had perhaps a more similar percentage of complaints compared to one another, and yet really didn't. So there was a proactive registrar that had so many more domain names, had nearly 14.5 million domain names, compared to a reactive registrar that had roughly three million domain names.

So while the reactive registrar managed 20 per cent as many domain names as the proactive registrar, they actually ended up having nearly

50 per cent as many complaints as the proactive registrar. So this was just a prime example of how being reactive really does not seem to be beneficial, because you're just inviting more abuse complaints. In addition to the analysis we were able to do to, to contrast the business costs and really see that there was a business incentive to be proactive, we asked our respondents for their suggestions on how to improve business costs associated with domain name abuse.

There seemed to be universally a call for improved DNS literacy on behalf of those lodging the complaints. Some of them suggested they wish people understood intellectual property law a bit more, and understood when not to complain about certain content related issues, whereas others really just wanted to see that abuse complaints were routed to the appropriate party. So if in fact that registrar was the responsible party, they were more than happy to get the abuse complaint, but if in fact it was an abuse complaint that should in stead go elsewhere, such as to a hosting company, they wanted to see that routed appropriately.

As everyone knows though, this can often be a gray area where perhaps it might be that multiple parties are the appropriate parties to receive abuse complaints. That was something where we saw that as perhaps an education opportunity, to improve literacy on at least domain name abuse, and anti-abuse processes, in hopes that that can make it easier for everyone around. Similarly, something that was suggested by one respondent was that perhaps there should be a model form that's used for abuse complaints.

Because if there was a model form that asked for specific information about domain name abuse complaints, the theory goes that then you would at least have enough relevant information on that form that you could cut down on the back and forth between the registrar and the complainant, or between the registrar and their customer, because you would have a bit more information from the get-go. So this is something that the SDF really hopes we can contribute to in perhaps working with the community and working with our partners to draft a model form.

The final summary of our report, we really showed that it's really more expensive to be reactive than proactive, because bad customers are bad for business. Moreover, being proactive seems to make registrars less attractive for would-be cyber criminals. If a cyber-criminal knows that there's a specific registrar that really takes proactive steps, doesn't just let anyone through the floodgates, then that's not going to be as attractive for cyber criminals to go to, and in fact would be good for good customers and the would-be bad customers will go elsewhere.

Going forward, we really hope that we can continue to do research like this down the road but get even more participation, and get participation from registries and hosting companies as well. We'd love to be able to get enough data to actually make some distinctions between different business models and look at how those that employ resellers versus those that do direct sales maybe draw a distinction in their anti-abuse efforts, and the business costs associated with it.

We'd also like to analyze technical data for correlations, and essentially identify really cost effective models that we can promote to the community that can be good for business and good for anti-abuse. If you would like to - and I strongly encourage you to - take a look at our report, it's on the website today, right on the homepage. Just go to securedomain.org. Feel free to reach out to me, or reach out to Norm Richie if you have any questions after this presentation, in case we don't have time right now. Thank you.

EBERHARD LISSE: We're running a little bit late, so I can take one question. Jay Daley?

JAY DALEY: Thanks. Jay Daley from .nz. You showed two different registrars - one that's proactive, one that's not - one that's complaints based, and the difference they got. Have you actually got a statistical analysis of those two different types, that shows whether that specific case is a generality?

DREW BAGLEY: As I mentioned, on the one hand we were able to get enough participants that represents 12 per cent of the domain name market, however we did not get as many participants as we had hoped to perhaps apply this to the whole market. This is what we saw with our data, with our respondents.

JAY DALEY: I don't mean with the whole market - that is a step. I mean within those that you did get…

DREW BAGLEY: Within those that we did get, yes, absolutely.

JAY DALEY: Is there a statistically significant difference?

DREW BAGLEY: There is a statistically significant finding that being reactive costs more money than being proactive.

JAY DALEY: Okay, and do you have a quantification of the difference in cost of that significance?

DREW BAGLEY: Yes, in the report we do.

JAY DALEY: Okay, thanks.

EBERHARD LISSE: Again, thank you very much. How it's to be applied to very small ccTLDs is another question, but it's good to look into this. Next one is Ed Lewis.

ED LEWIS: Good morning. I'm Ed Lewis. I work for ICANN staff. We are doing some things to the root zone's KSK, or planning to do in the near future, so I'm going to give a set of slides here to talk about what we're doing and what's happening as we go. My agenda is the usual one. I'm setting the scene, talking about the HSM stuff we're doing, talking about the KSK stuff we're doing. I'm using acronyms now because I'll explain them in line coming up, and then the big finish is a little less big than I planned it would be.

Some background. The root zone KSK - this is the trust anchor for the DNSSEC hierarchy throughout the DNS as we know it. We have had a root zone KSK in place since 2010. The key itself has not changed in all that time - in five years. Also, until a few months ago we hadn't even changed the HSMs, which are one of the important pieces of hardware in the system, and so we're about to make changes to both of that, so we'll talk about that today.

After five years of operation there is some concerns about the HSMs. There is also a contractual requirement to roll the KSK in some documents somewhere, so that's the starting motivation. The players involved with this, the root zone management partners consist of ICANN, consists of the US Department of Commerce NTIA and VeriSign. They're the three organizations that have been involved with the root zone management for many years. Also this year we have gathered an external design team to help us review the KSK roll plans.

ICANN is doing this work as the KSK manager, under the IANA functions contract. VeriSign is involved as the ZSK manager - just to give you an idea of the separation between VeriSign and ICANN in this instance. So what is a KSK? A KSK is a key-signing key. It is a public private key pair that is used to validate the very top set of keys in the DNS that are being used by those who validate their DNSSEC today. The public key of this KSK is distributed widely. Anyone doing validation of DNSSEC has to basically make a copy of this key.

There are automated ways to get it, but you have to actually copy this key and install it somehow - sometimes part of a distribution of software, sometimes by yourself - but this key has to be copied everywhere for this to work. This is why it's going to become a tricky operation. The private key is only used within the HSM. That's pretty easy to change. That's not the real issue here for the KSK roll.

The HSM stands for hardware security module. Many have probably heard about those. They are an important piece of a DNSSEC operation. It's not necessarily for the protocol, but it's important to operations in some areas. It's a specialized piece of hardware. It operates the KSK. The KSK is inside of it and no one actually ever sees the KSK in the raw anywhere. We actually give this device data to sign, and it gives us back the signatures.

Now, public impact of all this. I'm going to switch my order a little bit coming up. Actually, now. The HSM change is not something that should impact anybody, unless it goes wrong. HSMs are devices inside the network, they're a bump in the wire. As we change that out, by

plan there should be no visible impact of this, so it shouldn't be a problem, as long as nothing goes wrong. The concern right now is that the ones we have are getting old, but I'll say that that's just a concern - there is no scientific evidence that the age of them is causing us any issues at this point, so we're not at a critical or emergency point of view at this point.

The KSK roll is going to have a different, a larger impact, because anyone who's validating right now, if you don't follow along the change of the KSK, all of your DNSSEC will say no, and nothing will validate. It has to be copied out there, and that doesn't mean that people have to do the copying. There's a lot of work to do here to make sure we all know about this and so on, so that's part of what we're looking at as work to be done.

The presentation I have today is informing you of these updates, and also to call attention to an upcoming PCP that should open next week, if all goes to plan, on a document that's been under review and under preparation by the Design Team. We're putting it together this week. It's a draft report, and we'd like to get people to look at this. It's going to be technical. We'd like people who have experience with DNS to take a look at it, people who have concerns about the roll to look at it, and let us know if we're looking at everything we need to worry about? Worry is a good word. And do we look like we're on the right track to put this together?

Now, there are two means of feedback that I want to point out. One is the informal method, which is talk to one of us - talk to one of the

Design Team Members, who I haven't named yet and I will soon - using mailing lists. A lot of us are on lots of mailing lists, but this is all informal. The formal approach is the PCP, which is why that's important to ICANN's processes - that is the place that ICANN makes sure we have the response. We see comments, and we will acknowledge them that way. But either method is fine - just remember that if you want a formal way, come to the PCP, virtually.

HSM change, I'll say one slide about that. It's a straightforward replacement - we're taking the old one out and putting the new one in. It's a little more complicated than that in reality. We have done that in one of the two facilities. We had a ceremony to do that in April. That actually went flawlessly. It was one of our better-followed ceremonies in terms of script and actually what happened. The second ceremony, to do this at the West Coast Facility in California is going to come up in a month and a half, and the documentation for this plan is at the URL at the bottom of the screen, which says what the approach is to do this change. If you have questions about that, let me know.

The KSK role - compared to the HSM this is something much more significant; greater public impact. There are a lot more ways to do this. This has been building over the last couple of years. In 2012 ICANN had a public consultation about the approach to this. In 2013 there was an internal engineering effort to come up with a plan for this between the partners. This year, in 2015, we convened an external design team to take a look at all this work and tell us if we were on the right path, and whether we should go forward. Because frankly the technology is

changing so fast, we want to make sure we're making the right decisions.

The current plan for the team is to study through this month. Right now we're wrapping up a draft report. The report is going to be out for public comment for 40 days, the usual length of time, opening up right after this ICANN 53 Meeting. Then following that, the design team is going to come back and respond to their comments. Then the partners will come up with the plan for executing this and going forward. I see someone who looks like he wants to ask a question. Okay.

So Design Team [unclear 00:57:13]. These are the people on the Design Team. Joe Abley is here. He was at the mic earlier. [Yashiro] is here, I bet. If you're in the room, can you wave your hand? There he is back there. Jaap Akkerhuis is probably not in the room right now. Paul [Vauter] was in the room, his bag is in the room… Paul? Anyway, these are people you can contact in person here, along with me and other people involved with the root zone operations. These are the folks that are looking at this. The other folks were at other meetings, they're not here this time around.

In theory, on paper, the KSK has been rolled by many people, successfully and unsuccessfully. That's important because we know what works and what doesn't work, but in [unclear 00:58:07] case, we've built up a lot of experience in this. We know what works, what breaks. But the root zone is a little bit different, because in this case, for most of the TLDs that have done it, you just go to IANA and say, "I'm

making changes," and it works. In this case, there's no one above IANA to say, "Here's the DS, accept everybody. Everybody come and get my new key."

So we have a little bit of a different problem this time, even though it's been done before. There is a mitigation. There is an IETF document that says how to automate the distribution of these keys, and that should help. However we have a lot questions about this. Is it going to work? Will validators get all these responses? Will the automated trust anchor we plan to move out there actually work? Will operators know how to prepare, react? Who's going to get the brunt of the complaints? Will all of these code paths work?

This is part of what's been going on for the past year, is looking at these questions. I'm going to give you a preview here. This is the rough outline of the document being edited as we speak. This is mostly to give you an idea of what's going to be there and what you might want to talk about. If you have questions here, if we have time, one or two can be raised. I think we're running a little behind on the agenda though.

First, the history, scope and motivation for this is mentioned. Cryptographic considerations; looking at the algorithm and such, are we doing the right thing, is the bit length the right thing? Protocol considerations such as DNS is incredibly sensitive to the size of the response, and can we keep our responses small enough so that this critical piece of the Internet keeps working while we're making these

changes? Operational considerations - are all the players involved going to be able to manage all of this uncertainly at one time?

The impact on DNSSEC validation - there's a lot of study going into who's actually validating today. We've learnt quite a bit of stuff there, even external to this process. And how will the validators be able to follow along with the change in the key? Even if they don't change the key, how bad is it? Trust anchor publication - we're looking at trying to make sure that people know about that. There's concern about testing. We're looking at ways of putting some platforms out there that people can run their code against to see will I follow the plan?

We're also talking about the plan itself. There was also an engineering plan from 2013 of how to do this, and we've been making sure it's the best way to go. Then also an analysis of the risks - what could possibly go wrong and how are we going to anticipate that happening? We're trying to get that in the document. These are to give you an idea of what will be in the draft, coming up to the public comment period, and from my last slide I'm going to give you some URLs to go and look at for DNSSEC information if you want to look there. Go ahead Warren. I know you want to ask a question.

WARREN KUMARI: So back in 2013, SSAC published SSAC 63 and it had a bunch of findings. The most important one was that ICANN and a bunch of other people should immediately undertake a significant worldwide communication effort to publicize this. I don't know if anything's happened with that, but I haven't seen much communication.

ED LEWIS: Haven't is a past tense. Yes, that is part of what was noted in this review, that this is very important. In fact the interpersonal communication of this happening is very important, and we're looking at what was done in 2010 and whether or not what was done in 2010 was the right way to do this, or what are the other methods for this. Yes, this is all about getting people to realize that this is a change you've got to do a manual thing to.

WARREN KUMARI: Another thing I wanted to mention is myself and Rick Lamb each have test beds, so if people want to test their 5011 processing stuff, they can point [their BIND 01:02:09] or Rick's, which I think works a lot better, and then at least have an idea of 5011 implementations that work.

ED LEWIS: I'll say that your test bed is actually a really good asset. I've used it and tested it with BIND and Unbound, and I've managed to get them to work as much as I could. In fact, I think that's one of the things you'll see highlighted in the testing part, is that ahead of time we'll have test beds like that available. I've asked the distributors of the code , "Does it work with that test bed?" and so on. So that's actually a very important thing. It's an enabling element to this. Wes?

WES HARDAKER: How are you Ed? A couple of things. One, unfortunately I think most validators probably aren't doing 5011, so how much that's going to help us is questionable at this point. Now, whether they should be is

one of those things that should go in the ICANN's outreach methods. Two, I remember seeing some sort of measurement statistic in the past, but I can't recall it at all, about how long it takes a root.hints file update to propagate.

This is fundamentally a similar problem on the upside, because the validator community right now is on the leading edge they're probably following it a little closer than resolvers that have been out running forever and ever. Unfortunately, that still is probably a tale we need to address and look at as the worst case scenario that we're likely to hit up the road.

ED LEWIS: Yes, you're hitting on topics that have come up in discussions. First of all, there's getting software to even do this stuff. That's one. Two is to make sure those that carry that software around, distributors of OS's for example, to make sure they're packaging the right stuff, to make sure that's being conveyed out there, and then there's also some observations about the root validators out there that seem to be unmanaged, and what's that going to do? I won't get into that here. It's a philosophical and too long discussion for the room, but that's a good set of questions to ask. We want to make sure if it's not seen, put it in the public comment.

WES HARDAKER: On the upside, I think you hit an important point; that automatic updates by OS vendors are much better than they were a decade ago

where those types of files and configuration files were getting pushed out by OS vendors with much better accuracy than in the past. There's still a lot of people running old versions of operating systems, but we won't go there.

ED LEWIS: Right, and to answer what Warren said about public outreach, it's not just about speaking at conferences, it's about this…

EBERHARD LISSE: I'm going to cut this off now. Can you take this offline please? We're running a few minutes behind. I'd like to keep on schedule. Next is the host presentation from .ar. We've lost that one. Following that one would be Dave Soltero from ICANN DNS Engineering. Are you in the room? You can start. We'll take that one and see what we can do with the other presentation.

DAVE SOLTERO: Hello everyone. My name is David Soltero. I work for the ICANN DNS Engineering Group. We have just renamed ourselves from ICANN DNS Operations, as we feel the name fits better what we're really doing. My presentation is more about L root, what we're doing and what we're going to be rolling out in the near future. You guys know DNS can be a very complex thing that happens really fast, and also the users don't know about our work. At the first step we have those things called root servers that start making the DNS work, and ICANN runs one of them, L root, one of 13 root servers running on the AS2144.

We've got IPv4, IPv6 since 2008, and we're actually renumbering IPv6 addressing in the near future. [unclear 01:07:35] is based on two different configurations. We run three clusters. They're managed on an ICANN [caller] space or provided by a third party. We've got one in Los Angeles, one in Prague, and one in Western Virginia. In addition to that, we have another 147 instances that are hosted by partners. Today we run NSD, by NetLabs, and we're in a process of rolling out Node DNS from ccNIC in September.

There are pictures of what it looks like. I have three pictures in there. One is on clusters and one is one of these [u clear 01:08:39] we actually run. The one below there is a portable L root installation, which we're trying out actually in this meeting. The benefits for Anycast are known, but basically they're bringing the service closer to the user, lowering the RTT, improving the user experience, and also increasing capacity as well, and reducing the likelihood of certain types of [tax 01:09:19] infrastructure, and also having some management - we can bring nodes in and out of service for maintenance, without really having an impact on our users.

Here's a picture of where our ethos is today. We have 150 in service in 73 countries, and we'd like more. We've got all the common ways of identifying our Anycast instances, with using host name BIND, ID servers, NS-ID, and with IFC [privilege] by Joe Abley and Terry Manderson, we also have a [server called 01:10:05] identity.l.root-servers.org. That let's us publish information of each of our nodes.

For those of you who want to host an L root instance, we call it an [L singo 01:10:25]. The organization needs to be capable or willing to sign a non-disclosure agreement with us, purchase an appliance. This is new. We used to allow the organization to buy a server matching a set of specs, but that's now changed. The organization needs to provide us floor space power, and now we require v4 and v6 connectivity, and also be able to establish a BGP session with us. If the organization can satisfy the prerequisites then we sign an NDA.

They complete a technical sheet with IP information and routing information. We'll sign a contract. Once the documents are signed we will install the machine remotely. This process could take as little as a week, and some cases it takes a month or longer. It depends on how quickly the organization can proceed with providing technical information and signing contracts. Our latest instance is the one for Argentina, NIC Argentina. It went live in January 9th and we started receiving traffic from their ISPs some time on the 12th, at [unclear 01:12:12] per second there.

The graphic I just showed you comes from Hedgehog [unclear 01:12:19] software [unclear]. New [web n] for presenting the data. We have the information almost in real-time being published on that URL, hedgehog.dns.icann.org. Some of the features we provide. Also we leave the software as open license so anyone can get a copy of it. Everything you want to know about it we will file on the DNS.org domain. This graph of Hedgehog for D-17 showing [pathways] for the five regions; Asia Pacific, Europe, Africa, North America and Latin America.

Behind L root and the DNS Engineering Group is John Bond, myself, Mauricio Vergara and Terry Manderson, who's the Director. Goals - to have a DNS excellence, extension and diversification through the L root worldwide, collaboration with our peers, get the best engineering process and transparency and documented processes. How to engage with us? Well, we participate at DNS OARC and research bodies, on the [dogs 01:14:06], the IETF, and many of our mailing lists, and of course on our website and Twitter. That's our email. You can contact us on our website or via email, at DNS@icann.org. Any questions?

EBERHARD LISSE: I have a question. If I wanted to do this, on one of these links, where will I get these specifications? The requirements that we need to be able to fulfill to approach you?

DAVE SOLTERO: Well, I mentioned them before. The requirements are basically this. Be willing to sign an NDA and contract with us, it's a zero money contract, it's just to establish our SLA between the two parties. You should be willing to purchase an appliance, and we'll give you information on how to buy that appliance once you engage in the process. Be able to provide v4 v6 connectivity, hosting and power for that appliance, and be able to set up a BGP session with us.

EBERHARD LISSE: My point is, we as .sd registry, we don't have anything to do with addresses, so we have to talk to our host about this. That's why I need

a specification that I can give to them, about, "Can you set this up?" But we'll talk about it offline, I think. Any other questions? Okay, thank you very much. Next one will be Joe Abley.

JOE ABLEY: Good morning. We still have six minutes, so it's still true in this time zone. I have a few thoughts here. I think I'm going to give you some time back on the Agenda, depending on how many questions we get. I wanted to share some thinking that's been occupying our minds at Dyn, as we build on our infrastructure, and see what we can learn from the audience here. The story starts back years and years ago. There was a time, I believe, before I went there to work at ICANN, when L root was a desktop PC that had been rescued from the garbage and installed professionally on top of the fridge in the break room and that was L root.

Back in a certain length of time ago, that was perfectly fine, and lots of important servers were very modest when you saw them in person. I'm sure everybody here, whose beard is as at least as gray as mind, has a story about how the server would mysteriously turn itself off at 2AM every morning, and it turns out that's when the cleaner plugs in the vacuum cleaner. Anyway, we've come a long way since then, we think, because it used to be that that was fine, then that wasn't fine. Then we needed a more protected environment, and we built a machine room. A machine room was just a room with a locked door.

Then it was a room with a locked door and better air conditioning, because the servers were generating lots of heat, and then it was a

locked door, and air conditioning and fire suppression, because we suddenly realized if this room catches fire we're in big trouble, and then maybe we have raised floors, because raised floors make it easer to run cables, and it turned out we didn't know how many cables we needed to start with, and we had to think a bit harder about that. Then we need cable trays at the top, and then we need cabinets and we need all kinds of stuff, and suddenly we've found ourselves with a really big investment.

It's expensive to run, it's certainly expensive to build. But whenever you have something that's very expensive to build yourself, you know there's somebody else that's going to build it for you, and charge you per month to do it, and then we have data centers. Suddenly the economics are different. It doesn't make any sense any more to spend thousands and tens of thousands of dollars up front to build your own machine room, when you can just go and rent a cage at Equinix or Switch and Data.

So suddenly that's the new thing - you don't build it yourself, you don't worry about the power consumption and the cooling and the access control, you buy all these services from somebody else and they compete with each other, and they do it at scale so they're cheaper. Suddenly we had machine on fridge, then we had machine room, and now we have data center. Then of course the density of these servers continues to increase and the cooling requirements at scale become quite ridiculous and we start having to worry about liquid cooled servers and running all kinds of ridiculous things around racks that weren't imagined when data servers were first built.

Then it turns out that maybe it's easier just to get somebody else to deal with all of this, and not have to worry about what the guy in the next cage is doing with the hot aisle and the cold aisle, and whether he's leaving huge mounds of cardboard in the aisle, which is causing a fire hazard, and all these kinds of stuff. Forget about the physical servers - let's just use the cloud, because the cloud will solve all of our problems. This punch-line has been delayed. There it is.

All together, "Cloud!" The cloud is fantastic. The cloud makes all these things easy. All the things that were difficult and required logistics, people and hardware, and things breaking and people falling over cables disappear fro our lives, and we just have simple rest APIs, and we turned all of our hardware problems into software problems, which is good, because fundamentally we're all software people at heart, and hardware is just an annoyance we don't want to deal with, which is a great story. It's certainly better than the machine on the fridge.

But we still have to remember that this cloud in the sky is not some sort of artificial domain that is ephemeral and everywhere at all times and things like that. This is not religion - even if that cloud looked like religion. We have to remember where these things are. If we have services that were building for our office, if the cloud is close to our office then we're good, it's fine. But if we're trying to build services in a cloud and we want to deliver them to everybody, everywhere, then we have to remember that the cloud isn't everywhere. For lots of people the cloud is data center in Ashburn, Virginia, which is great if you live near Ashburn, Virginia. But if you live far away across oceans, it's not so great.

In particular, it's not great when the stuff you're trying to send to people, to end users, is large content - video or software updates, stuff like that. Suddenly, distributing this stuff centrally doesn't look like such a good idea, because there are limits to how fast you can push the stuff, and there are limits to how big these pipes that together comprise the inter-tubes can carry. There is only so much data we can force down them. So we have our centralized resources, which are good for all those reasons, but bad for these reasons. What comes next?

Maybe this chronologically is not correct, but to augment this content that we have in the middle, we have content delivery networks. So for the content specifically we can push that out as close to the people who are going to consume it as possible, so that they can get it locally. We still put all our compute in the cloud, all the webpages and everything else like that that build the user experience and manage the user experience, all still happen in places where it's just managed by software, and then we push the content, which is the really expensive thing to get to the user, right out to the edge.

So then we have these content delivery networks. We have Akamai, we have other people who do hybrids of these things. We have Google, we have all kinds of people building stuff out - Netflix, for those that have Netflix in their country, another fantastic example. The idea is that you do all your user management in the middle, and when you're finally directing them to some content, or in the case of Netflix you say, "It's not available for live-streaming," in my experience.

But if it is, if you're watching the same thing as everybody else in your town then you can get it from the local cache, which means the only bandwidth that gets congested is between your ISP and you. It's all very predictable, and you don't end up with problems like… As it turns out you do end up with problems like IOS7 or IOS8 - the Apple phone update that almost brought the Internet to its knees. Maybe you don't have it quite so much. Maybe people learned from those sorts of things, and I think that's probably true.

Then we've got a hybrid situation then. This is the modern way to design your service without having to touch any bits of hardware yourself. You've got cloud to provide all the applications and all the service management and logic, the compute, and then you've got content distributed through a CDM, and you have a choice of cloud providers. You don't just have to trust one, you can use multiples, and you have a choice of CDM providers. Some have different coverage from others, they serve different kinds of audience, so you have options, and you can use multiples together.

Companies like Dyn have emerged to fill this little gap in the middle. If you are the service provider, how do you know which CDM to direct your user to? Maybe it's the one closest to the user. Maybe it's the one where I want to keep down less than a certain amount of volume per month because it costs me less. Maybe I provide some arbitrage services between CDNs and things like that. In any case, there's this optimization layer that Dyn has found a way of filling, and it's working very nicely for Dyn.

Of course, Dyn's service here, filling that gap between managing the end user experience, itself, it's a service. Itself needs to live somewhere, and where does that live? Does that live in the core? Does it live right on the edge? Is it content? It's not really either of those things, it's kind of both. The way that Dyn is helping people deliver these end user experiences or these excellent end user experiences - I'm not a marketing person or salesperson, so I'm not familiar with these elevated pitches - but these are the kinds of words I hear.

We built out this infrastructure. We have an Anycast infrastructure for TLDs and for other infrastructure zones. We have a measurement platform that collects trace, rout and ping data as well as PGP data. We have archives that go back a decade. What else do we do? We provide what the purest would probably call "disgusting DNS tricks" at the edge, where the response that we give is not the same for everyone. It depends on who the client is, it depends on the time of day. It depends on all kinds of weird things that make any purist in the DNS Working Group cringe. We do all that nasty stuff.

We need to have a platform that has that computational element but also has that location. We need computation with location. So we don't need a cloud. What we need is a swarm. Here you can't even see the sky. This is not in the sky, it's a swarm, it's in your shirt, it's stinging you in your armpits. If you're afraid of birds, these are not birds. These are insects, and if you're afraid of insects, don't worry, these are birds. I couldn't really tell what they were - they were just a stock image from Google.

In any case, that's our idea. Instead of having the cloud, which is far away and ephemeral, this thing is really close to the end users, but also brings the computation, not just the content. I've described some of these things before in more speculative presentations going back a couple of years, but we actually have those now, and we're building it out, and it works. I'm here to talk about it. We have these compute resources that are very widely distributed, that are designed for massive scale. In some previous conversations about this people have said, "How many nodes are enough?" and I've said, "How many AS's are there, 50,000? Let's start there then."

Let's not assume that we're only going to go for one per continent or 20 globally or something else like that. Let's assume that we can go into every network where it makes sense. We'll go multiple times into the same network if it's a big network. It doesn't matter to us. We can manage as many boxes as people want to deploy. We've also noticed that other people that build out close to the end user have made the decision to optimize for a particular type of application. They're optimized for serving video, or they're optimized for serving static content or something. A lot of those people probably regret that decision, so we built something that's far more general purpose.

We can ship out docker images to the edge, and orchestrate them, bring them up and have them communicate with each other in a very flexible way. The way I talk about this internally to the developers is you're driving into work in the morning, and you have an idea while you're in the car. You take out your laptop, and there's your persistence layer, and there's your data collection, and you throw

together a docker image that's based on all five bits that are already written, plus your particular five lines of Ruby that you think needs to go out.

You press a button and it's distributed across 300 nodes. You collet your data and it's live in the Internet, and it's an Anycast service. By 11:30 you think, "Yeah, I think I've had enough now," so you close it down, and then you have lunch. We just deployed a brand new service in two hours because we felt like it. We didn't necessarily have to tell anybody we were doing it. There was no harm or risk to any existing service. We have a generic platform where the resources are dedicated to the things that are important but also flexible enough to be able to throw any application we like out there. We think this is good.

It's not envisioned right now as a public cloud service, but we could always work with partners to develop a docker image that made sense for them, that could run on our platform as well. It's flexible. That's the point of it. As far as the ISP is concerned, what does it look like? It looks like a 1u Dell. It looks very familiar. It ships from the factory, you plug it in, you stick a USB key in it, if it doesn't already come with an appropriate image burned on at the factory. You turn it on, give it some addresses through the console, serial or VGA, and there you go, it's live.

The rest of the testing can happen remotely. Services can migrate there. You, as a network partner, can say, "I'm in Texas. I don't want gambling.com, because that's going to make me very upset. So don't put that on here, but these other things I'll have, that's fine." Build the

constraints around, and then solve for the global set of constraints that says what is allowed to run where, where are the resources? Where do we need the capacity? Then just push these things out.

As I note here, it's a value driven partnership with the ISPs. Just like your Netflix, and your Google cache and all these kinds of things, there's benefit to the ISP in hosting this stuff, because you no longer have to trust five other networks to reach and reserve Twitter.com. You can do it locally. Here's my sales slide, which you've been pre-warned about, because it's not page by page. I mentioned I'm not the salesperson. I was told I was allowed to have a sales slide. We need a booth babe. Matt Schumann, can you stand up please?

When you do want to buy something, that's the man you go over and say, "No, that's too expensive," and he looks sad, so you have to buy him a drink. That's the way it works. That's the process. Matt is the person here from sales. The thing we have running on here today is Dyn TLD, which is why I thought this was a good audience. We have lots of people here involved in running TLDs of various flavors. This is a fairly new platform. We've got 55 TLDs live. It's a mixture of ccTLDs and gTLDs.

At the moment they're all running on BIND9, but they could be running on anything. Each TLD customer runs in their own docker container, so if you want to use Knot, or you want to use NSD or something like that, you can use any of those things. We'll work with you to make that work. We collect new gTLD compliance stats. For people who like DSC, lots of cc's like DSC, so we decided, "Let's use exactly the stock,

standard, Duane Wessels approved DSC collector, and not try and improve it or anything like that." Is the data exactly like Duane intended? Yes it is. That was the decision there.

Zone analysis by Dyn research - you can read all this stuff. The interesting thing I think is that we can work with you and say, "Where do you want servers?" Do you want to say, as you might to other providers, "I want my TLD hosted in each of the continents, apart from Antarctica"? Or do you want to say, "Actually, I think we've got a particular community in this country and we'd really like to have local service in that network"? We can satisfy those kinds of requests very easily.

Give us an introduction, just as Dave was talking about with L root. Give us an introduction, we'll get a box shipped, and we can turn up a new node for you in really as much time as it takes Fedex to deliver the box. This is the area of flexibility. This is location plus computation. That is the end of my presentation.

EBERHARD LISSE: Any questions?

JOE ABLEY: All right, so if anybody is interested in the sales side and wants to buy service, or can't afford service but wants it for free anyway - that's making Matt very sad now - Matt Schumann is back there. Thank you very much for listening.

EBERHARD LISSE: Thank you very much. We're a little ahead of time. Next would be Joao Damas. I'm still missing Luciano from .ar. If he doesn't pitch up by lunchtime, we will take one of our presentations and substitute it.

JOAO DAMAS: Hi, I'm Joao Damas. I'm working for [unclear 01:33:00] and doing some research and development on the DNS. I'd like to start by thanking Eberhard Lisse for having me here, and also DNS.pt, the .pt ccTLD for making it possible for me to be here this time. This talk is a little bit different in the sense that in the TLD we look at traffic that comes into our authoritative servers, and this is always traffic that is generated by recursors. It's very rare to be able to see directly traffic generated by end users, because they usually stand behind recursive resolvers.

I was curious about what that traffic might look like, and I started looking at it. This is what I'm talking about. You are all familiar with this type of picture. Normally you see the right hand side of this graph, the traffic generated between the recursive servers, and the NS authoritatives, particularly if you run them. There's a whole other set of traffic that you never see, and particularly the recursive resolvers have an important role in that they say there's a lot of stuff that gets to be seen at the authoritative servers, but they usually have a cache.

That significantly filters out the amount of traffic you see on the right hand side compared to the left hand side, because if you ask a

question that someone else in your network has asked recently, the recursive resolver would probably be able to answer from the cache, instead of going through the whole process of asking again. The end customers run software that's quite different from what typically is used for recursive resolvers, and so it has different properties. I wanted to look at that as well.

The problem with looking at this sort of data is that it has direct information about who the end users are, in the sense that you see the individual IP addresses of each of the host machines at people's homes, that generate the traffic, and that quickly generates a lot of concerns from people that don't want that information to be exposed. So it's a bit tricky to get your hands on this sort of data. But now there are automization tools, you just tell people who have some trust, and eventually I was able to get this data.

Looking now, this is why this is a [unclear 01:35:40] more recent set of data, because for different reasons this set of data I'm using here is maybe one euro, maybe a little bit more. Things change all the time on the Internet, for good or for bad, and I'm trying to get this updated. One of the questions that Eberhard had for me was… I'll tell you a little bit about the software that is used to do this sort of analysis.

All the heavy lifting and narrowing down of interesting things in the big data samples - and there was a lot of data - was a sample that was 24 hours of DNS data at two decent sized ISPs of recursive resolvers, so there was quite a lot of data. There is a good tool out there called PacketQ that is very good for getting a sense of where things are going.

This is a tool developed originally by IAS, .sc TLD operator. It's available on Github. They unfortunately decided to not continue maintaining it one or two years ago.

I know there is some maintenance going on inside of NetNode, but they are not yet making that new code available. Hopefully that will change in the future. In any case, the tool does work, and since the DNS hasn't changed dramatically in the last two years or so, since they stopped maintaining it, the tool is perfectly capable of looking at everything you want. It lacks a couple of features, and that's why there is some additional software that maybe I'll make available, once I'm not embarrassed by showing the source to everyone. Basically, 90 per cent of what you need to look at your code…

For your data, I would say this also holds true for authoritative data. You can look at it with PacketQ. It's very nice and surprisingly fast. What this thing does basically is it uses pcat as if they were MySQL databases, without the need for MySQL. It gives you an interface where you write SQL-like queries, and then it transforms that into lookups into the pcat file. It's astonishingly fast. If you ever have an issue with trying to find out how your server is seeing the traffic, or what people are doing to your server, this is a very handy tool to be aware of.

It does take a little bit of time to get used to it and make the most of it, but if you know a bit of SQL it's quite easy to get your head around it. The source is available on Github. Let's look at the data. This is what the samples looked like throughout the whole 24 hours. 4,000 to 5,000,

maybe a bit bigger at sometimes, queries per second, through a 24-hour period. Then I'm going to go first through some basic DNS data to characterize the sample. Luckily, as one would expect from a running service, most of the answers you see are no other answers, so the DNS works.

Surprisingly, there is a very constant baseline of queries that get refused, or get returned not implemented, which probably points to software doing weird stuff at the users' sides. If I get the new data, I'll look more into that. If you look at what domains people are querying, it's also quite interesting to see the patterns. This is one of those cases where you will not see this traffic at your authoritative servers, because the cache is going to dampen what gets out.

There are two peaks there. They are way above everything else, basically, and I hope ISPs, if they care, are talking to the [unclear 01:40:06] those, because apparently these are NETGEAR devices that use always the same time, every day, to go check what's going on. So there is this massive onslaught of DNS queries at NETGEAR.com in every 24-hour period. It seems a bit odd, to say the least. Just to be able to see these things better, this is the domains-related infrastructure. You can clearly see the NETGEAR peaks there, which are kind of annoying if you're running an ISP.

If you take those out of the graph then you start seeing everything else. You can see where people are getting their content from, which of the CDNs are being used in a given world. You see Akamai, you see the quite irrelevant iCloud there, and everything in-between, so Yahoo,

Google and everything. You see who is serving ads towards you, and this can be useful. Some of the CDNs provide appliances that you can install in your network, so if you know which CDNs are being used by users, you can then go and talk to the CDNs and ask, "Can you put an appliance here?" and know why you're doing it and what the benefit is.

If you look at entertaining sites, just out of curiosity, the big red graph there is YouTube, and the huge peak there is what a viral video looks like in terms of DNS. It's high and narrow. People move on very quickly to whatever is next in their lives. Then there's a bunch of stuff that shouldn't really be there. Those are the .local queries that escape all the networks, the wpad.home and assorted queries like that. Not much you can do about it, except coping with it. Other interesting services are security services. The PDF for this is uploaded, so I'll let you go there if you want.

In the social networks, Facebook is apparently the king, and then you get to brand websites, and Google is apparently where everyone goes to do most of their work these days. Going into what capabilities users have, and whether it makes a difference between the traffic that an ISP would see, and what an authoritative server would see, then we look at things like EDNS usage. There are current studies by APNIC, [John] Michaels and Jeff Huston, that are seeing an increasing number of EDNS-enabled packets going towards authoritative servers, to around 40-50 per cent, something like that.

That is not actually what you see inside the ISPs' networks, where basically the usage of EDNS from things like laptops or any Windows

machine is actually negligible. The little trickle of traffic that is EDNS-enabled that you see down there is basically due to people who have their own resolvers inside their networks - so small businesses that hook up to their ISP using their service facilities, or geeky people who… The end laptops, the end operating systems, still do not make use of that.

This is also visible from the people who use EDNS, how many are carrying the DO flag, the DNSSEC okay flag in their queries? That's actually most of them. That again points to the fact that all these EDNS queries are actually generated by people who decide to install their own resolvers; small businesses typically. But again, this is similar to what you see on a authoritative site. If you look at those that set the DO flag, so they say they will be willing to receive DNSSEC data, how many, when you send them the data, actually come back and check anything with it? That's one way of looking at this; is looking at the DNSSEC key searches, then that is very little.

It's only the hardcore geeks amongst the geeks that actually look at this. It seems that at the end user level, DNSSEC so far has very little visibility. No one is using it, and that's basically because neither OSX, neither Microsoft is actually performing that kind of verification at the very edge. One other thing we wanted to look at is the TTL of the responses. So far, the DNS has been able to scale very nicely because of caches, and the fact that you can ask off a TLD or another authoritative serves a question and then use that response multiple times towards your end users makes the authoritative server of DNS do very large numbers of users.

The problem with traditional values for TTLs, which used to be set around one day, typically, is that they limit how fast you can change things. In these days of clouds and shifting attentions and viral videos and all these different behaviors that you see from what used to be the traditional use of the Internet, people want to have changes that they place much quicker than one day. Also, throughout history, even people who knew what they were doing have made mistakes, and if you have a big TTL you are stuck with your mistake until every cache times out, typically in a day.

So people are basically moving to lower and lower TTLs, and this has an impact on the efficiency of caches. I wanted to start looking at this problem by looking at what the solution is of TTLs. Also, some discussions that happen even at the IETF, that go sometimes to inform development and further standards work, sometimes people that discuss these things are still in the mind frame that most TTLs have at least one day long. What this graph is showing is that that is no longer true. Unlike in the other previous graphs, where the Y axis was a linear access, in this one the Y axis is logarithmic.

What this is telling you is that people are moving massively to very, very short TTLs. You still see the peaks at the default values of one day, half-day, two days, three days, but if you look at the graph you'll see those at the 10-1,000 query level, whereas the smaller ones, the range is five minutes are at the million. That's something that I think is worth looking more into and seeing what can be done, and how this affects actual work at the ISPs, and whether authoritative mechanisms

like pushing content from authoritative servers closer to recursive servers can have an advantage.

Of course, you need to be sure that this is properly done, and that you don't just move the propagation problem from one side to the other, but are not actually solving it. That's stuff I plan to look at, particularly if I get my hands on a more updated data set. That's what I have for today. If you have any questions about the data, or the software and how you use it, I'm very open to answer them.

EBERHARD LISSE: Two things. We've never looked into this in detail, but we have a saying in Namibia, that on Sunday morning you have the best through-put, because that's when the guys, the people who always look at the porn sites, are in church. I've used the same, PacketQ. I used DNS-CAP to look at the data on the wire. It generates a binary format, which is difficult to interpret for the uneducated user like me, but PacketQ allows you to apply SQL language to that, which is very easy, and then I figured out it's a relatively simple way to load it into a text file and then load it into MySQL in big batches. That works extremely well.

I've done this now for almost a yea rand stopped it after a hundred million records, because the analysis even in MySQL just takes too long when I want to make nice graphics. It takes about four to five hours to generate a picture. But we have also used Owl as the statistical analysis software, which can access the MySQL database directly. There is surely other databases and longitudinal ones and whatever, with which you can do this on a larger scale, but even to look at this

was a relatively simple method. It's quite easy to do, and it's even very efficient. We also find that it loads extremely fast. Any other questions.

Thank you very much. We are a little bit ahead in time. It appears that we are losing our host presentation due to translation issues; that we are unable to arrange for a Spanish/English translator, due to reasons beyond our control. Jay has asked for some time to speak about the service level agreements, or expectations that the CWG Stewardship is developing with IANA, or asking IANA to develop. I think this is the right audience for him to say a few words about it. Have you got a presentation?

JAY DALEY: Before the slides come up, my name is Jay Daley from .na. I am one of five registry people who are from a Working Group called Design Team A, who are looking at developing something strange called a service level expectation for post-transition IANA. We may be more used to the phrase service level agreement. I'm not entirely sure what the difference is, but we're working on an SLE, not SLA. I'm only going to go briefly through this. I've got half an hour but won't be using anything like that. Don't panic. We'll be going to lunch a little early.

There are three gTLDs on this and three ccTLDs. The ccTLDs are me from .nz, Patricio, who is somewhere in the room, and Paul Kane, who's not a ccNSO Member but has been around for many years. The three gTLD representatives are Elaine Pruis from Donuts, Jeffrey Eckhaus from Rightside and Jeff Neuman from a variety of places. We

have looked at the initial SLA that currently exists for IANA as agreed with the NTIA, and it is our view that it is entirely inadequate for the way that we want to go forward. It is inadequate in a number of respects.

Firstly, the level of measurement that takes place is considerably less than we believe is necessary. There is not the full service definition that many of us would know. It is not broken down into the different stages, and those things are not measured. This has some problems. There is a perception among many people that change requests to IANA are going through a two-tier process where contracted TLDs are going to get done quite quickly, and non-contracted TLDs are going to the back of the queue, but they're all being done within the overall time allowed.

Now, we want to disprove that quite thoroughly by having a proper measurement that shows for each stage of the process how long things have taken. There is another concern, which is that IANA is deliberately holding, or is responsible for holding up some changes, and we want a measurement that stops the clock when IANA transfers responsibility back to the requestor, say by asking for clarification, and then starts the clock when it becomes IANA's job again. We want to see that level of detail move forward, and that level of detail is simply not available in the current SLA.

The second thing is that the current SLA is very generous in certain respects currently. We were asked to look at the breakdown of services, the IANA functions operations, to look at the existing SLA, and

produce an updated capture of those processes and then develop a new SLE. That's what we were working on. We believe that as customers of IANA we need to state what we think is the minimally acceptable service level that we require, what reporting requirements we have, and what breached levels.

If you take something such as name server changes, we might say that a name server change should be processed within two days 80 per cent of the time. So we set a target of two days, and we then set a breach level, which gives IANA some headroom above the two days, by saying 80 per cent, and that reflects the priority that we put onto that. If we thought it was a very high priority then it might be two days but 95 per cent breach threshold. So they need to achieve it 95 per cent of the time.

We are not recommending any changes to the process at all. It's not our job to try to force automation of IANA through this process. That's IANA's job to do that. All we're trying to do is follow their processes. We are though recommending a major change to their IT systems by requiring this amount of measurement and extraction of data to take place. Now, for .nz, the registry that I run, has 650,000 domain names. Monthly we produce a 20-page report to our own internal regulator, just on registry performance.

We also produce another 15-page report on marketing performance. We do this type of measurement quite normally, quite naturally. I think for many of you in the room, that's the same level of reporting detail that you go into. We are not aiming to push IANA anywhere

beyond that - in fact I don't think we're even going to reach that level that I've described that we do, but we are going to take it from the current level, much closer to that.

We have some principles. I'll be quick about these. We want attributable measures. We want to measure who is responsible for something and how long they take to do something, because we want to be clear what IANA is responsible for and what they're not responsible for. We want to measure, as well as individual parts of a process, the overall process, so that we can understand trends. We only want to collect things that are relevant. This is important, because whenever you discuss data collection, people have very good ideas and they want to know all sorts of things that are not relevant.

We don't want to burden IANA with that, so this is just about relevance. We want to clearly define what those are, so that everybody understands where we're going with those. We want the thresholds that I've mentioned already for breached levels. We want a review process so that this can continually approve, and we want regular reporting so that people can see where things are. Now, we have gone through a process that has had a lot of work to take place.

Initially we weren't sure how much ICANN would engage in this. Our concern initially was that there would be a push from ICANN Senior Executives for us to just stick with the current NTIA SLA, but that's been replaced with full engagement from IANA, who are now working hard to help us map out the processes and negotiate and agree what levels

of performance are required. That's going well. Some of this has moved on a bit. I'm not going to worry too much about those.

I'm going to show you a few slides of data, which I haven't produced. But we do have considerable data available on changes, based on historical data, how long things have taken, why they've taken certain lengths of time, and that's been used to set an initial set of performance targets, which we're negotiating with IANA, and an initial set of breach thresholds, so that we can understand that a certain percentage have taken place within a particular time, and others haven't worked.

On the slide deck that's available for download with part of this meeting, there are numerous slides about this, which go into a great deal of detail on stats and the raw data is available I believe as well. That's effectively it. There isn't more to get into so far. You are as much customers as we are customers. If there's any feedback then please pass it onto Patricio and I. If you've got any specific questions, hopefully they will have a document out relatively soon with our initial targets and things in there as agreed with IANA.

But generally, I think that now that a lot of the work is out of the way, this is a week where we're going to be trying to communicate the work we've done. I'm here talking today and will be talking tomorrow at the CcNSO Day as well. Please approach us if there are any questions, and we'll take it from there.

EBERHARD LISSE: Thank you very much. I don't know whether you all noticed - Kim Davies is going to speak after lunch about IANA workflow. This presentation was short notice, I'd rather have had it straight before IANA, with Kim in the room, for a little bit better engagement. I've given Kim permission to use my last request to IANA for a WHOIS change, so there is no privacy issue. I requested a change of street address and a change of name server service, because one of our secondary providers changed their platform.

The software was clever enough to find out that it's two changes, and asked me to split it in two. The [street 02:03:10] change took five days, with which I don't have a problem. The name server change took five days until the software sent me an email that it needs manual intervention, and then it turns out to be some problem with some EPP communication with VeriSign to actually implement it, because when the human intervention was made it took about ten seconds to finalize the old request. I'm not arguing it must be two days.

I fully agree with what Jay and others say. We must have quite a bit predictability. We want to know how long this is going to take, and then we want to see how many times, or how many per cent of the times, it's taking that long. Personally, I feel we should hold IANA to the same measurement standards that ICANN requests from their gTLDs. They request a really detailed, specific measurements on a monthly basis, but are basically not willing to provide the same level of measurements or the same amount of measurements for whatever reasons that I'm not really aware of. Anybody else?

Okay, so then I'll close the proceedings a little bit early. We are going to meet at 14:00. Kim Davies will speak soon. Then Andrew Sullivan will speak about the CARIS Workshop. Jacques Latour will present, NIC .cr will present. Adiel Akplogan will give us an introduction into the ICANN technical engagement to which he has been appointed. Then so from 15:40 until 17:00 we will have a PGP signing party.

In case you don't know what that is, many of us used PGP or GPG keys, but the keys in themselves only show that the person who sends the key has access to the key, but it doesn't prove that he is actually the person he is presumed to be. For this, human intervention is required, which needs a certain process, which we will go through. We had a key ring advertised on what is called Biglumber. It's a website set up for [unclear]. We have about 22 participants submitting their keys. If all of them are here, we'll get them all done.

The ones that are not here, or the ones who want to load up their keys now, we have set up on every day of the week until Thursday, at 13:00, one hour at [unclear] A, B, or C - it's listed somewhere - we can then sign the [streglos 02:06:03]. Not all will be able to be there to sign, but the people who come, they can sign each other, so we at least can get something going. The hope would be that we would eventually do this on a regular basis so it becomes part of a regular event. The more people who sign each other's keys, the more trustful the keys are. That's about it. Let's go to lunch. Self-catering this time again. [applause]

[tech-2-22jun15-en]

EBERHARD LISSE: Okay, good afternoon. Can everybody please settle down and sit down? Obviously the people who are present need to be punished for the ones who didn't come. It's the usual thing, but this is after lunch - it's always a problem. Kim has asked for some time to explain our workflow. As I said earlier, I've given him permission to use my last name server and WHOIS change, without any respect to my privacy. He can disclose anything that's concerned, if he wants to, in case there is an example that is needed.

KIM DAVIES: Thank you very much, Eberhard. Thank you for the opportunity to talk to you today. What I wanted to focus on is our current thinking with respect to how we do technical checks to the root zone. Those that aren't familiar, there are a series of technical checks that we do. I'm going to walk you through them. It's been almost ten years since we last looked at them and re-evaluated how we do things. I think in ten years a lot has changed. The industry is very different, there are a lot more players. Technical configurations have evolved.

It seems like the right time to sit down, re-evaluate, take stock of what we're doing, and perhaps work out a roadmap for where we want to go. With that in mind, here are the basics. Whenever you submit a change request for a zone that we maintain - and I'm focusing on the root zone, but we also apply these for .int and some of the other domains that we manage - we conduct a series of technical checks. These technical checks are designed to check for errors, or other issues

that should require further analysis before a change request can be implemented.

The second point is we repeat these at several intervals throughout the life of a change request. When you submit change requests to the root zone we do a technical check then near the beginning. We also do it again just prior to implementation in the root zone. Then as a third step, VeriSign, as part of their process for implementing change requests will do a third set of technical checks at that time as well. In the life of a change request you'll probably get tested three times. All tests are fully automated.

I think this is an important property. Everyone wants more automation with respect to our business processes, and therefore having a set of tests that can be fully automated is pretty essential to that. Now, if you're a customer and you submit a change request and you fail it or you don't pass it, what we do is we send you a transcript, we ask you to remedy any of the issues that we've identified. Absent to any response, our system will automatically retest.

If there is a transient issue and it just resolves itself automatically, like maybe one of our requests to you just didn't get through, and it works shortly thereafter, it will automatically retest, it will automatically pass, and it will automatically proceed without anyone intervening. That being said, as a customer you have the ability to go in and do retests, so if you fix something you can force it to refresh faster than the retest. The other point I'd make here is that the process has built in

a process where customers can appeal against technical check validation failure.

You, as a customer of ours, can say, "I know it's failing it for these reasons. Here are the mitigations in place. We think you should proceed." That will go to one of our subject matter experts, usually me, and we'll review it and make an assessment and we'll talk to you about that. Some of the reasons why you might want to make an appeal I'll cover as I go through.

Just as a background, the current technical checks that we use today were set as a result of public consultation in 2006. We did a public consultation. A bunch of ccTLD and gTLD registries responded back then. It's all posted there. As I mentioned, it's codified in our systesm and our tools. Let me walk you relatively quickly through the current tests that we do today. These are the really basic tests. I don't need to elaborate too much - minimum of two name servers that don't share an IP address. You need to have valid host names that comply with RFC 1123, and your servers must answer authoritatively.

In terms of network connectivity tests, you must respond over both UDP and TCP. We need to be able to send a UDP and TCP query to each of your name servers, and they need to respond back to us. Network diversity - they must be in two topologically separate networks, defined as not sharing the same origin AS. We assess this through inspection of routing tables. We use third-party services like RIPE, [unclear 00:10:11] and so on to do that, just because they have extensive routing monitoring networks in place already. You can't use

tunnels, you can't use private networks. I think private networks would probably fail other tests anyway, but it's there for completeness.

Consistency - there are three key checks we do for consistency in terms of the data in your zone and in our zone. Any glue that you want to put in the root zone must match the authoritative records for those host names. We checked the A and the AAAA records in the authoritative zone for your name servers and make sure they match what the glue is you want to put in the root. There must be consistency between the delegation and the zone. So the NS set you want to list in the root zone must match the NS set in the apex of your TLD. All pretty basic DNS stuff.

Then we check for consistency between the different authoritative name servers themselves. We define this as having matching NS sets at the apex and matching SOA sets at the apex. When we say SOA, what that really means is that the serial numbers match, in effect. Other tests we do - we're checking that referrals from the root do not truncate. If you have too many name servers listed, or they have a sub-optimal naming scheme, then if you send a maximum sized query to the root servers, the root servers will truncate the response.

This is where the number 13 for the number of root servers derives from, is this 512-byte limit of packet sizes in non-EDNS zero contexts, and we do the calculation on the fly for submissions, and you need to fit within that 512-byte limit. We also check for open requests of name service. You cannot provide open requests of name service on the

same server that is authoritative for your domain. In terms of DNSSEC, we checked for DS record format, just normal stuff there.

We checked to make sure that for each DS record there's a matching DNS key at the apex of the zone, and we check for validation, that we can validate the [RI sig 00:12:30] by using that DS record set that you've provided us. That's the list. Now, let me tell you about some of the things we've run into, particularly more recently as new gTLD vendors have started rolling out a lot of TLDs. A lot of these discussions are becoming more and more frequent.

In terms of network diversity, when the test was specified in 2006, it was a vast improvement over the way we did it before then. There was a network diversity test, but it was all to do with allocations and prefixes of IP addresses, and now we have some visibility into the routing table that gives us some kind of more scientific, objective way of coming to the conclusion regarding network diversity.

The assumption back then, when we were having discussions in the community was that all good TLD operators, the best current practice is that you engage some other company to at least run one of your name servers. They'll host it somewhere else and they'll do that for you. I think that assumption went into the way that the test was architected. I think what we're seeing now is an evolution in the industry where, for good practice reasons, some vendors want to maintain all the infrastructure themselves. I think there's a discussion to be had about the pros and cons of that, and how that might apply to the future evolution of IANA's checks.

Now, one piece of feedback we often get as an appeal is simple, "It's all in the same AS, but I use Anycast, so don't worry about it." I think it's very important to consider that we're not talking about failure of a particular network location. Diversity is beyond simply that. We're talking about things like someone has screwed up some configuration or software roll out across an organization. It could be business failure at a higher level. Maybe the company simply goes out of business or has something radical; a court injunction, something like that, that impacts the broader scale of the company.

If we accept that TLDs are a critical Internet infrastructure, the configuration should be such that the name server configuration should be able to sustain any kinds of risks like that. I think it's important to be mindful, when we're talking about diversity, to think about other kinds of risks, not just the risks that a particular location goes offline. What we've also seen is some vendors have obtained a second AS, essentially just to pass our test. Now, that was never the intention. That's a false sense of diversity.

We don't want to get into the business of putting in arbitrary rules that people just skirt around. If that is in practice what's happening, then that's something we should address. We want something that's practical, but achieves the job. Maybe that's wishful thinking. Maybe there's no perfect way to define it, but I think that's our goal. I think importantly we need to recognize there are unskillful operators out there running TLD - perhaps none in this room, by virtue of the fact you're here.

Some TLDs are run by people that don't understand the DNS very well. That's a reality. It's still a reality. It's less of a problem today than it was some years back, but it's still a problem. Be mindful that IANA's job is to have tests that work for everyone, not just experts. So we need to keep that in mind. Some other issues. This is still a vast minority, not very common, but some TLDs wish to list inactive, standby DS records in our zone. Where this breaks down is a design assumption that any technical data you put in the root zone we can somehow cross-verify.

If you put a DS record in, we can match it to a DNS key. If you put the NS records in we can match it to your NS set. There's a built-in integrity check in the DNS that allows us to do those checks. We don't have that with standby keys. So I think it's worth recognizing the demand is there, but we don't have a technical check accommodation for that scenario, so that's something to be mindful of. Just an example - this is broken, but I think it's a god example - occasionally we have TLDs submitting DS records that don't have the SEB bit set. Essentially they're giving us the ZSKs instead of the KSKs.

Now, these validate fine. There's nothing technically wrong with it, per se, but it's a pretty good indicator that there's something wrong. You would think it's a no-brainer - we should add a new test, check for that, reject them. Well, the times it's happened we've gone back to the customer and they've said, "No-no, go ahead. We want them to be that way." So it's not always clear-cut as exactly how we should be treating these. Again, the topic worthy of discussion. Serial coherency

- the way it's normally meant to work is the first line, is that we check NS1, NS2, NS3 - all the serial numbers match up, that passes the test.

In the rare event that we do a test just as the zone is updated, it's that middle line. Maybe we check NS1 and it happens to have the next version of the zone file. NS2, NS3 have the older version. That's fine. It will fail the tests at that moment, but we do a re-test, and the theory is it's like that third line, where upon retest everything is coherent again, we get a matching SOA and we can proceed. That's fine, and that still works in 99 per cent of cases, but there are some zones like this.

They're loosely coherent, they're never in-synch with one another at any moment in time; think about zones that change every few seconds and they have a cluster of name servers it gets pushed out too. There are some smart ways to deal with this, but the current test, as defined right now, cannot deal with this scenario without manual intervention. Some other feedback we get - I picked a few quotes off some DNS operations mailing lists I'm sure all of you are on. "ICANN should be testing and blocking TLDs until these network blocks are removed." "We have ICANN checking query rates and uptimes but not protocol basics prior to letting TLDs go live."

"ICANN and TLDs should be showing leadership in this area." I just show them as examples that we do get feedback from non-direct customers that say IANA should be upping the bar and testing for more things. I'm not saying it's right or wrong, but it is a piece of feedback that we get from the community. Another piece of feedback is that right now we have stricter requirements for v4 versus v6. It's an open

question. Maybe times have evolved where we should change the v6 and v4 requirements to be the same in terms of how IANA implements them.

What can we do? Re-emphasizing what I said earlier, I think the biggest takeaway is all checks need to accommodate all TLDs regardless of skill level, and also add to that they need to be automatable. I started off by saying this - our goal is to do PCPs similar to the way we did in 2006. Some of the specific ideas to consider is what I just mentioned - some of those things like serial number coherence, network diversity, standby DS keys and so on. You probably have other ideas as well. I'm very happy to hear them. They're some of the ones that are top of mind in terms of what we hear back from our customers.

Another idea to consider is technical check wavers - the idea that for certain types of checks where we might agree that there should be a permanent ability to remove that check for a given TLD. Take that serial number example. If you are a TLD that is constantly having that problem, maybe we just disable that check for your TLD - some kind of opt-in process. You need to acknowledge the risks and so on, but I'm sure we could work that out.

Another piece of feedback we get is that the technical check feedback is not very intuitive - it's hard to understand. I fully agree. It's not very good, the way it gives you feedback, so we're re-implementing it from scratch. The system will have comprehensive de-bugging logs. You, as system administrators, will be able to go in and get those diagnostic

logs out of [unclear 00:21:54] self-service. You can look at the trace routes and pings and all that kind of stuff, dig results from our system. I think that will help a lot in diagnosing some of the issues people see.

Right now, if they see a cryptic error message they have to email us and we have to have a dialogue. I think a lot of the issues that we see, most people are fully equipped to solve them themselves, if they have the right information. IANA can regularly start performing the checks for TLDs, just as an ongoing basis. Right now we only do them when you do a change request. We could just start doing them regularly and send you a courtesy email notifying you, "Hey, we've noticed this new condition we didn't notice before. Just a heads-up. can we opt-in? Opt-out? It doesn't really matter to me."

Also, for certain kinds of changes, like if your NS set suddenly deviates from the one we have, maybe we could even send you a pre-populated change form and say, "Hey, we noticed you added an NS record. Do you want to make the change in the root zone? If so, click this link, and it's all pre-populated for you." So just some ideas about how we can be a bit more clever about this than the way we're doing it today.

We could also provide tools so you can do the checking yourself. Either we can give you our code, or… I was at a meeting a few weeks ago with the center project zone check, .fr and .se have written tools, and in that tool there's a profile where you can choose, "Do the IANA tests." So making sure that works the way our tests work, so you have a tool available to you to do them on your own. Finally, just some other ideas - CDS support; another way of triggering change requests perhaps. I

mentioned in the beginning we do tests multiple times to our change request. As the lifetime of change requests get shorter and shorter…

That was designed for the days when it could take weeks to do a change request, because in those weeks your configuration might have drifted and more issues might arise. If change requests are being done very quickly, like within a day, maybe we don't need to do it three times, so maybe we work out a new strategy for that as well. I think that's about it. That's just a quick run-down. Good ideas are welcome to me at any time. I'm working on a public comment period. I'm not sure exactly when, but in the next few months I hope we can get that posted.

Some of those technical things I mentioned, like making our error messages more understandable and so on, that's work already underway right now that we're working on.

EBERHARD LISSE: All right. Thank you very much. Jay first.

JAY DALEY: Should there be more than just a PCP? Should there be a better mechanism for community involvement in this that sees maybe some of these ideas originating within a community group and then coming that way?

KIM DAVIES: It's a good question. I haven't considered that. There are pros and cons to both. I think either way we want buy-in, so whether it's community driven or whether the buy-in is through consultation through this and getting comments on the record and so on doesn't matter to me.

JAY DALEY: All right, because I think it's good that you come forward with things - that's obvious, because you're seeing the sharp edge of it - but I suspect that some of the things for example about two vendors, I'm pretty hard-lined about that. You have to have two vendors. As well as running [my own 00:25:23], I have two separate vendors as well, so I actually have three. But it would be useful to test the mood of that in a different way and see if there are new things that come out, that's all.

KIM DAVIES: I agree that IANA is in a sensitive position. We can't force standards on people, so whatever we're testing for has to be the will of the community, let's say.

JACQUES LATOUR: Jacques, with .ca. It's about the emails that you send for changes. There's a link in there. You just click on it and it approves a change right off the bat, without authentication or anything. Is that part of what you're looking at to get rid of and just use web portal? Or what's the strategy around that?

KIM DAVIES: The email links for authorization has a nonce in it that is unique to each email. So root zone authentication has…

JACQUES LATOUR: Sent clear text everywhere.

KIM DAVIES: That is true. I presented at a previous meeting, but our goal there is to introduce a whole new authorization model, having talked about it in this presentation. The idea is with this new authorization model we'll have two-factor authentication, and importantly opt-in - you can opt-in and say, "Never approve change requests for my TLD based on an email. Someone always has to log in with the user name and password, second-factor, before you accept a change request."

There's a lot of complexity in migrating from the current model to the new model and having a transitional phase that's making it hard, but that's a work in progress. That's the way we're trying to address that.

ROBERT MARTIN-LEGÉNE: Hi, I'm Robert from PCH. Thank you for changing the error messages. We get many questions about that. I have only a request that when you do your test, if you can do it on a clear cache somehow, because sometimes people hated something like the IP address is incorrect of one of the glue records, so they go and change it and then BAM!

They're hit by a two-day TTL or something and the have to wait. That seems to be what's happening sometimes.

KIM DAVIES: We were seeing that. It's actually a bug. It should have been fixed about a month ago. We handed off recursive lookups in the testing suite to BIND and then BIND was tuned to not cache everything, but it wasn't working the way we expected it too. We just stripped that out and changed the approach relatively recently. If you're still seeing it, please let me know.

ROBERT: This has been there for years. I haven't checked recently, but…

KIM DAVIES: Let me know if you see it again.

EBERHARD LISSE: Okay, no more questions. Thank you very much for coming. Next presenter is Andrew Sullivan. I haven't seen him yet. He can come forwards please? Have you got a haircut, or what?

ANDREW SULLIVAN: And a tie!

EBERHARD LISSE: Don't worry, we won't cut it off.

ANDREW SULLIVAN: I'm Andrew, and I'm here to talk a little bit about this Workshop, which we ran in Berlin last Friday. It was coordinating a TAC response at Internet scale, and in case you don't know, I am the Chair of the IAB but I'm not speaking for the IAB today. Normally they don't let me wear ties, so this is a really fun adventure. I'm not speaking for the IAB, I just happened to be there because I'm on the IAB. I went to the workshop, and these are some impressions from me.

Normally when we produce workshops we have a report afterwards, and it comes out as an Internet draft, and then you have the chance to kick it and spit on it and so on. So this is an opportunity that will be coming to you, so if you want to know what the people in the Workshop thought then that Workshop Report Draft will come out and you can send comments and say, "Why didn't you talk about this?" Either the people working on it will say, "Oh yes, we did talk about that."

Then they'll add something to the report, or else they'll say, "Gee, that's a good question. We should have invited you to the Workshop, and we'll do that next time." So don't send comments if you don't want to do more work in the future. Apart from that, comments are valuable. You'll see that on the IETF list. It will show up anytime soon. I'm hoping it will be very soon. Kathleen Moriarty, who is one of the Area Directors for Security at the IETF, she was one of the main organizers and she's very diligent in the way that she does things, so I expect that the Report will come out early.

Now that I've told you all of these things about why you should take my point of view with a grain of salt, the question is what is this thing? We had a workshop, and this was part of the FIRST conference last week. FIRST is… I can't remember the expansion. You know emergency response teams for Internet stuff? This is the meeting of all the people who do that kind of thing. They met in Berlin last week, so we held our meeting on the end of it, and the point of this is to have a Workshop of interested people.

Many people don't know what the Internet Architecture Board does, and that in fact includes many of the people on the Internet Architecture Board. But whatever we do, we're supposed to be interested in paying attention to the overall architecture of the Internet - how the various parts fit together, and because I have this time, and you have to listen to me because I'm talking, I'm going to tell you a thing or two about how I think architecture on the Internet is important.

There are different views about what this word "architecture" means, but there's a great earlier history of architecture, building architecture, that involved a Roman guy by the name of Vitruvius, and he wrote the so-called Ten Books on Architecture and so on. This is an earlier tradition. It's not like Corbusier and those kinds of people who think that it's all about art. Vitruvius thought that what made for beautiful architecture was that things be fitted to their purpose - this was a major focus for him.

For instance, to Vitruvius it would be ugly to put a marble building in the middle of a desert, because there's not a lot of marble in deserts, so you have to bring it from far away and it would be complicated and so on, and if you've got sandstorms in the area then that tends to wear down the marble, and so it's not a good material for that environment. These are the kinds of considerations that are important in Vitruvian architecture - the fitness to purpose is really important. So I think that architecture on the Internet, one of the key things you have to pay attention to is how do things fit together, given what you're trying to do? We're trying to hook networks together, and interoperation is one of the most important things.

Therefore I believe that incident response on the Internet is actually one of our great architectural weaknesses. It would be like building an office building today and failing to think about putting in conduit or allowing for RF between floors or something like that, so you couldn't use Wi-Fi. Maybe 40 years ago it would have been fine not to worry about network cabling in your building - in fact, I worked in such a building.

It was wired very carefully for all the electricity you'd ever need for typewriters, and the only problem was when you wanted to run a network cable you had to call the cord-drilling company and they'd come and run their cord-drill all night so they could drill a hole through the eight-foot thick concrete walls. It was a real pleasure to wire that building. So I think that this is an architectural point that we're missing on the Internet today - that our ability to respond to

emergencies on the Internet actually depends on people knowing each other.

You know a few people and you know who to call, and so you get into trouble and you call them up, and then they can fix problems. That's a way a lot of incident response on the Internet works. This tendency to rely on informal relationships is both a strength and a weakness. It's a big strength, and the reason I think it's a big strength is because if you have that sort of informal relationship then you've built some trust with the person, so all of the trust that you'd need to do…

I don't know how many of you have had the pleasure of noticing any of the IANA transition stuff that's going on, and all of the accountability things, and yes, I am wearing a tie - but in there yesterday there was some discussion about trust and institutional trust. What people want to do of course is build up that mechanism. You notice when mechanisms for trust have to be built up, instead of, "I'm going to go to the bar and have a drink with this person and we'll exchange PGP keys, and now we can trust each other enough to respond to a network emergency," that's the way we've traditionally done it in the operational community.

In the more formal mechanisms, what you do is you set up a committee, and it has an accountability measure of another committee, and that committee is appointed by all the people in the universe, carefully divided and with a modulus of three - and then you get this incredibly complicated system, and people still don't understand how it works, and then they ask, "What about my point of

view? How come my point of view isn't represented here?" However, this, "How come my point of view isn't represented here?" is also a problem with these informal relationships from the other direction.

That is, if you want to go and explain to your Ministry of Telecommunications of whatever, "How do we respond to emergencies on the Internet?" it is not a satisfying answer to them to say, "Well, ten of us got together and exchanged PGP keys in the bar, and now we trust each other - don't worry about it, we'll look after any routing emergency." That doesn't make people happy, that makes them feel uncomfortable, and I can see why. We've got this tension, and this is the reason that we thought it was important to run this Workshop - to see if we can do something about this.

There is another really important feature to this. Historically, operators have gotten together in rooms like this, and presumably in bars like downstairs. They've known each other, they've exchanged some information, so that when something goes wrong you contact one another, call each other up, send an email. However, what you don't have there is ease of joining. If some new operator comes along, they've got to get an introduction.

For somebody who's fairly outgoing and already speaks one of the languages that's in use or something like that, that may not be that hard, but for somebody who isn't at all these meetings or doesn't know who to talk to, or anything like that, it's hard to join the club. I'm sure that everyone here had that experience at one point, where you had to be introduced or something like that, or you're just an outsider and

you were standing on the edges. Well, as the Internet grows, it's growing right now, we're going to have a lot of people joining the operational community, because we can't all do it ourselves. Well, maybe all of you can, but I can't, and I don't want to stay up that much.

I want to make sure that one of the things that we have are mechanisms by which it's possible to join these communities and that should be relatively painless. It maybe doesn't mean that you can join in the deep way right away, but maybe you can get some ways in. What did we learn? Well, one thing I thought we learned is that reputation is still a really big problem. A couple of years ago there was a Working Group at the IETF that was intending to build a protocol to indicate reputation.

That is you know some people, and you'd have some way of indicating, "I know this person, so you can trust them this much," or something like that. This was pretty much a complete failure. It worked very well to document, in an elaborate way, a simple protocol that everyone is already using, which is the email black hole lists, but it didn't actually achieve very much else. It was a great disappointment to me. This turns out to be a problem for lots of people. A common complaint that we heard in the meeting was, "I have no way of knowing whether I should trust this other guy, and I don't have any way of getting started."

Another thing we noticed was that automation within particular communities is quite good. People settle on an interchanged format and they just rely on it, and you start to see that interchanged format,

you see somebody else producing it, and you say, "I can probably incorporate that into my feed." But if you've got different kinds of communities, for instance the D-DOS community and the TLD or DNS communities, and they haven't already agreed to all of this, then automation is almost impossible. People disagree about the formats, and you get engineers in a room, and what are you going to do?

Are you going to talk about what you should measure, or how you should measure it? "How you should measure it, for sure! Don't want to know about what you're going to measure, you want to know what's the format first. That's of course what we need to settle." This is played out over and over again among these communities so that there's a significant problem in these operations communities, that they've got formats that subsume other formats, and a bunch of people that say that CSVs ought to be good enough for everyone, and that consumes a lot of time.

There's really deep disagreement about what to share - so what kinds of attacks people are seeing and so on. This is still black art information shared among people who are in the know, and people who are working on another thing, who might turn out to be amplifiers, for instance, are very surprised by the effects others are seeing. Remember, these were mostly people who were all at the same Incident Response Conference that week, and they still were surprised by one another's results.

I think we see that here at ICANN. We certainly se it at the IETF, so I was shocked, but I was not surprised. I was shocked that people were

having this difficulty, but when I thought about it, it didn't surprise me very much. I did learn that FIRST has, like every formal group… I should say, IETF workshops are usually a bridgehead kind of thing. We hold the Workshop in order to learn something, but we don't know where we're going to go next. It usually doesn't mean we're trying to set up another formal group. I think in this case we're not trying to set up a formal group, were just trying to create some new communication pathways.

FIRST is apparently going to set up a special interest group to help on this, and [Marika Kao], who many of you will probably know, she is involved in that, so that may be somebody you might want to watch your emails for, for calls for participation and so on, because there may be something you can do there. There is another thing I learned that was very interested, and that was the amount of money being wasted on Internet response. The meeting happened as most workshops do, under Chatham House Rules, so I can't tell you who said this or what the context was, but I can tell you that it was reported.

They saved $500,000 a year by joining this one feed for $2,500, because they had five full-time staff people going around, chasing the problems they were having. This was infections stuff, and they managed to automate that. They replaced five full-time people with a simple shell script, and that's the kind of thing that is real money. That was $500,000 a year that they could free up for things they were trying to do, rather than fighting fires. You might want to look at this. What is it costing you in your operations to respond to help-desk requests about abuse? That may be a significant amount of money, and it seemed to

me that for this community that's an area that's probably worth looking at, and trying to figure out what could we do collectively to try to put things down there.

There's also a big disconnect among the policy people, the technical people and the lawyers. These three groups of people all think that the most important thing is a different thing than everybody else. A really interesting problem here of course is sharing information with competitors. I know about an attack that's happening. My colleagues in another company are under a similar attack, and yet the lawyers on both companies are like, "Whoa, we can't talk to our competitors, That's sharing competitive information."

The only way we're ever going to find these kinds of attacks is if we share that data. So as an industry, somehow we've got to figure out how to make that sane, and make policy and legal people in particular comfortable with that kind of data sharing. We're never going to solve these problems without it. The other thing about this - and this is maybe something we need to communicate better to both legal and technical folk - is that the data is going to get shared one way or another.

The question is whether it's going to get shared in a way that's good for you, or whether your attackers are going to share it amongst themselves, and they're going to keep that secret from you, and you're going to learn about data breach problems you have not from your competitors, but from the New York Times. That seems like an issue you want to worry about as well. The other good news - and I thought

I'd point this out here in particular - is that APWG was very, very kind in its remarks about its relationship with ICANN and the registry and registrar community.

They're happy, they think they're having big successes there, so that's actually an achievement where some data sharing, which has been not that expensive, as far as I understand, for anybody involved, has yielded some big benefits. They're able to shut down malware people, criminals, very much more quickly than they used to be able to do. Now, it's not clear to me - and I don't think it was clear to anybody else, but I can't be sure - whether that's going to scale given the large number of new TLDs. We'll see.

You've got a bunch of people - as Kim was just suggesting - some of the people who are coming into this community right now are not experienced with this. They maybe don't have the background and they might think, "We have to re-do the whole education of the lawyers, that this is new data, competitive advantage, blah blah blah," and it will be six months before you figure out, "Oh no, that's abuse data and we need to do something about it."

There also may be room, it turns out, for some more collaboration between names and numbers. It isn't clear to what extent we're using the intelligence we get about misbehaving IP addresses or IP address ranges, and misbehaving DNS names. Those two things are complementary data sets, yet I got the very strong impression that that data is not being shared very effectively. Those are all the things I had

to report about this meeting. I don't know if anybody has any questions?

EBERHARD LISSE: I was quite interested in this. Any questions?

[CHRISTIAN ESSELMAN]: Hi, I'm [Christian Esselman]. I'm with .nl, the registry for the Netherlands. I'm also chairing a Working Group called Secure Email Communications for Internet Response, and we actually have very much echoed what you're saying here. We have set up a basic incident response facility for ccTLDs, that enables them to quickly look up their contact information in case of large-scale incidents. We currently have around 134 subscribers, so 134 ccTLDs on that list, which is roughly 54 per cent of every ccTLD on the planet.

What I found very interesting about what you said was we also followed a similar approach. Our model is based on personal trust, which means everyone that's on the list needs to be on it with their personal information, like names, phone numbers and email addresses, and also we try to keep the barrier for joining as low as possible. It's different than what you're presenting here, but perhaps it's first go for more advanced incident response architecture.

ANDREW SULLIVAN: I like this. The thing that's a challenge about those kinds of environments are two-fold. I've been on some of these kinds of lists.

My job changes, because I work at a company that's in massive expansion mode, so my job title has changed on average every six months over the past three years. I got a new job, and I'm on maybe 900,000 lists, it sometimes seems, and I forget to update those things or pass the token on to somebody else. Then I'm on vacation and an incident comes along. One of the big questions we have is suppose this thing has to scale, suppose the number of ccTLDs expanded by an order of magnitude. What would we do about that?

It's unsatisfying to say, "We'll just use a roll email address," and you don't want that, because in the company that could be going to 10,000 people, and five of them could be actual scammers or something. That was one of the things we wrestled with quite a bit, and I'd be interested to see…

[CHRISTIAN ESSELMAN]: Scaling is obviously an issue - not so much in the ccTLD community, but if you expand to the g's, if you want to include the g's as well, then obviously there's a scaling issue.

EBERHARD LISSE: Last question from Warren Kumari.

WARREN KUMARI: What you were saying in the last response made me a little concerned. I happen to trust you, but I don't trust Dyn, and I happen to trust you and you might introduce Joe to me, but trust isn't transitive, no matter

how much you want it to be, and so I'm not going to trust Joe - and I've met Joe, so I have good reasons for that. Trying to scale this in any dimension without it being built on personal relationships feels like you're going to end up with a two-tier system - one of people I know and trust, and then another tier with lower amounts of trust.

ANDREW SULLIVAN: Yes, and that's part of the question. One of the questions is whether that's good enough. It could be that what you need for the barest sorts of stuff is this low-level, institutional trust that we can explain also within legal context and so on, and then say we've got that other layer, and that really is built on personal trust, and the law doesn't have any problem with the idea of personal trust. But as an institutional fact, it's going to be very hard to go to the world's governments and say, "No, actually, the world's routing system and all the TLDs, they all operate on handshakes by 12 people that we know. Don't worry, we've got it."

That's not an answer, and I think we're struggling to recognize that, because of course for years it has been an answer, it has worked pretty well. Thank you very much.

EBERHARD LISSE: Thank you very much for coming. Next will be Jacques Latour, and the following one is Maurice Oviedo from NIC .cr. I haven't seen him yet. There we go. The topic I have on the Agenda is slightly different from what he's presenting. That has something to do with I've only got a certain amount of space per line.

JACQUES LATOUR: I'm Jacques Latour with .ca. today I'm going to talk about a new model for provisioning DNSSEC, and that's to get the DNS operator involved in that model. We spent a lot of time at .ca to build our DNSSEC solution, to sign, to do all that stuff, to talk to ISPs, to get DNSSEC out in the wild in Canada, and the biggest challenge we have right now is getting signed delegations. Right now we have 104 signed domains out of 2.4 million. It's pretty sad. We've had our registrar sessions - we meet with our registrar a couple of times a year - and they're basically not interested in doing anything related to DNSSEC.

Every time they say the word "DNSSEC" it costs them money, and it doesn't generate value to them. So I started to look at that, and then as part our IETF in the [unclear 00:54:24] we had the DNSSEC DS auto [boff 00:54:29] and we talked about different ways of doing DNSSEC registration, or automating DS provisioning. This presentation came out after three or four months of digesting this and getting feedback from the community. The key thing is everything we're doing right now is based on the RRR model, and DNSSEC is all about the DNS operator. They're not included in that model. So we need to look at the model and figure out why it's like that.

So history, legacy and [sacred 00:55:08]. History, Nlnet Labs in 2004, they said the RRR model would not work with DNSSEC, and we need to look at something different to make it work, because the provisioning is not going to work. That was in 2004. The issue needed to be looked at, but nothing happened then. Then this is an email written by some

French guy somewhere. It says that, "Conceptaully, DNSSEC was developed - all the protocol, all the provisioning - based on the RRR model, and nobody stopped and figured out who should manage DNSSEC material. Is it the DNS operator?

"But that didn't happen, and then we had RFC 5910, which is EPP, and it's basically the provisioning of DNSSEC through the RRR model." That process is broken. Out of the 144 delegation we have, they were all done manually. None of our registrars use our EPP interface for .ca, not yet. 15 out of 150 registrars support DNSSEC through the WebGUI, but none of them use EPP, so that proves not a really good investment. If you look at the model, typically you register a domain, it goes to the registrar, they take care of all the registrant information and the name server information.

That gets published in the zone, there's a delegation. Then the interesting part is that the registrant can be the DNS operator for that zone, or the registrar, if that's part of their business. Or the registrant can get the hosting company to manage their website, and they manage a domain, and that hosting company can also outsource the DNS operation to a third-party, and their DNS operator. If we had DNSSEC in there, they're the ones that would sign the zone and create material that needs to make it up to the registrant, through the registrar, to the registry, to go in the zone.

It's broken. There's no way a DNS operator like CloudFlare or Akamai and Dyn, these guys, could get the DS from any of the zones they manage, they could have a different hosting provider in the back,

different registrars. It's super complicated, it doesn't work. The idea is to change the model - so to have the DNS operator talk directly to the registry. It's not the registrant talking to the registry, it's the DNS operator. So in this model, one, the legacy interface. The RRR model, you can still support that to submit DS and all that, but the idea is that the DNS operator would connect directly to the registry and initiate the bootstrap sequence.

That's to create the initial chain of trust for DNSSEC, and also ongoing with the maintenance interface. That means once it's signed, technically if you turn DNSSEC auto maintain on, we'll poll for the new DS and then the maintenance should not break in that model. More automation on the bootstrap, more automation on the maintenance, ongoing management of signed domains, and key rollover and automation.

So the process, this would have been through many, many iterations of this bootstrap process. This is when you create the initial chain of trust. You get a domain, you sign it. You need to securely get the DS to the registry to put it in the zone. We started looking at this, and this is the most simplest process we have. The DNS operator needs to prove they control the zone, and they need to prove they operate the zone. Basically, all the information you need to bootstrap a domain when it's signed, it's already in the zone. There's a DNS key.

So what we need is the operator to prove they control the zone. They add a text record with the key IDs of the DNS key that they need to be included to bootstrap. You sign it with DNSSEC, then you get a text

record that matches. That proves you control that zone and you want it signed. Operate is you prove that by submitting a request to the registry, like .ca, that you actually want to sign your zone, and the zone is ready. We talked about doing pass phrases, passwords and challenges and a bunch of stuff, but it doesn't really matter, because all you need there is a trigger to say, "Please sign me."

If everybody says, "Well, what if it's compromised?" Well, if something is compromised, you're compromised, it doesn't matter. Here it's super simple - you control your zone, you put a text record to say, "I want it signed," and then you click a trigger, web interface. I'll show a sample of the web interface. Once you say you want your domain signed, the verification process: the script is going to go in over TCP. They check all the signature for the domain if it's valid, so it's properly signed with the DNS key that's there, that the name servers match in the parent and the child, and that the text record actually matches the DNS key in the zone to be signed.

Basically it's a script that goes in, reads all the information, checks everything is good. It grabs the DNSSEC and key and basically generates a DS. That's about it. If somebody goes in and says, "I want to sign my zone," and it's already signed, you ignore the request, it's a duplicate. If somebody says, "Sign my zone," and it's not signed, it won't do anything. It's going to tell you it's not signed and why it didn't work. So the web interface cannot be used to attack a domain. Somebody needs to control the domain and say, "I want it signed."

Something like this: we built a prototype of this. Basically, the idea is that we have a web interface and an API for a registrar who [unclear 01:03:22]. The idea is they go to [Web], they say, "I want my domain signed," there's a validation engine, and then once it's all valid and it's good to be signed, then there's an EPP module to insert the DS for that zone inside the registry. That's about it. When I talked to a couple of registrars in Canada, to automate this, the registrar, the hosting portion of their registrar, they have to write code to integrate with their registry business to get the keys down. In this case, all they do is submit the name of the domain they want to sign. That's all they transfer.

So we have the actual code to play with. Paul [Vauta 01:03:32] is working with me on this to write scripts, so we have [restful] API and web interface. Yes, it needs a little bit of security and control. The prototype, we just did that last week or the week before. Control in maybe a registered email address of the person that wants to submit the request, to make sure that it's not anybody. This is a principle. You've got the web interface, you go here, all you put is the domain name, and go. That's all you need to do.

The domain is signed already, so go you go "Start". Everything is over TCP, due delegation. You check to make sure it's signed, you check the name server, you check the signature, you check the DNS key, is everything good? Is there a CDS? That's for fun. Then we calculate our own DS with the parameters that we want, and then we include that [3PP 01:04:34]. That piece is not built yet, but that's where it finishes.

Then we calculate the DS record and that would be included in the zone file in our registry, our database.

Input is domain name, click "Go". Go. I'm sure we can add more security around it. [Restful] API is similar - same thing, same result, just xml file. Do you have any questions? Then we started to look at the maintenance approach - ongoing maintenance using CDS. In the beginning I thought of going in and polling for CDS and then taking the DS format and including that in the zone automatically, ongoing, and I still need to figure out more about that. I'm not sure CDS is the right thing to do in this concept. Now we're going to write the registry. We're going to have a program that's going to pull all the signed zones, and we know which ones they are.

The maintenance can be as simple as getting the DNS key, and if there's two then we add two DS, and if there's one, we add one. We just mirror as is the DNS key, ongoing, and no need for CDS. Maybe you have a text record that says "new key" with the key ID, or something simpler than CDS. Then I can create my own DS. The DNS operator doesn't have to worry bout creating a record for me, for a registry. I create my own DS, I put it in the zone. I take a key, create the DS, put it in. We got it to work, but I'm not sure it's the right thing to do automated maintenance, but we're going to look at it, so feedback is…

Strategy. Right now I'm going to continue working on this. I need feedback. I think it's going somewhere. We have 104 signed delegations. None of our registrars want to do DNSSEC. Every time they see me, they run away, because they know I'm going to talk to

them about DNSSEC, so it's simply impacting my credentials there. We're going to make it work with some of our partners. We're going to write code, we're going to build it, we're going to put it in production, and then we're going to learn from that.

The code is going to be all open-source, so we're going to put that on the Internet somewhere. Then eventually once it works and we've got the bugs figured out, then we'll do new RFCs if we need to, or a best practice document or something. At the CENTR Meeting in Stockholm, Marco Davis came up with an observation that was pretty interesting. He said that if you disassociate DNSSEC registration from domain registration, it creates a registry lock.

Then if somebody hijacks a domain, they change a name server, they can't change DNSSEC keys, so instead of routing a domain somewhere else, it's just going to fail. So there's potential for enhancing security with this protocol concept. That's a lot of stuff. That was my last… My goal is to make the Internet a better place.

EBERHARD LISSE: You're probably the only one, but thank you very much. Robert first?

ROBERT MARTIN-LEGÉNE: Hello. Robert from PCH. I liked most of it actually. Have you considered what to do when people want to get out of DNSSEC? Because it's easy to subscribe people, but if somebody wants to, for some reason, un-sign, go unsigned for everything, do they need to…

JACQUES LATOUR: I haven't figured out that piece yet. There are multiple options. Today it's through the registry, EPP, or web - they delete the DS from the registry. So existing channel, I guess it works today.

ROBERT MARTIN-LEGÉNE: Yes, but that won't help the hoster, right?

JACQUES LATOUR: No, then you can go _undelegate text record and put your key… Yes. If it's simple to bootstrap, it should be more complicated to unsign, but…

WARREN KUMARI: One of the things that CDS gives you, which I don't think the other solution does, is it allows you to pre-publish DS records for keys you're not planning on using yet, which a bunch of people seem to believe is really important, and I don't entirely understand why, but they're really wedded to they'd like to pre-publish a DS record and then swap in the key when they need it, in an emergency.

SPEAKER: Well, for the first key you obviously have to have it done, you have to have it signed, because otherwise you won't [unclear 01:10:29]. So the first DS will actually always be the key you're actively using. It's not pre-published, so you cannot use this interface when you're not signed

yet. So afterwards you could, with the CDS record, update your original set, as long as you don't change the trust. So you're not allowed to go from signed to unsigned in some automated way. You have to have some kind of human securely say, "I'm going to go in secure."

As a note on the previous comment as well, the problem of doing an undelegate record or something is if you've lost your private key and that's why you want to go insecure, you cannot add it to your zone anymore. You can't sign your zone because you lost your key. I don't know how to fix that.

EBERHARD LISSE: I'll take two more questions.

PAT KANE: Pat Kane from VeriSign. Jacques, it's very interesting you've had these conversations with registrars. I'm assuming they're happy that you'd do this for them? So with the diagram that you had, the DNS operator, are they working on behalf of the registrant or on behalf of the registrar?

JACQUES LATOUR: The DNS operator can be anybody, right. It's whoever operates a DNS function. Your question is…?

PAT KANE: It's going to come down to who's going to be compensated, and by whom? The DNS operator is going to compensate for this. Do you envision it be the registrar, or the registrant, in the conversations you've had to-date?

JACQUES LATOUR: Compensated…

PAT KANE: Who's going to pay who?

JACQUES LATOUR: It's free. For us it's free.

PAT KANE: It's interesting, because one of the things we've been thinking about in this model is not just the DS record, but do you envision unbundling other resource records so that you're putting more domain management into the hands of the registrant and basically unbundling the registrar?

JACQUES LATOUR: Crawl, walk, run. We know we're crawling.

PAT KANE: Okay, because I'm sitting in the back going, "This is really interesting." Because you're unbundling the registrar and becoming pro-

competitive for the registrant to be able to provide services. We had the same issue at VeriSign, in terms of getting people to adopt DNSSEC.

JACQUES LATOUR: Ideally, you want the DNS… One idea I had was to, somewhere in the registry, to allow a registrar… Say GoDaddy enabled CloudFlare to manage the host and the DNS. That was a concept that we had. The DNS operator can manage, if there's a D-DOS and they want to move 10,000 domains right away, then the DNS operator could modify the registry on their behalf, to be more agile to responding to the D-DOS. I'm just focusing on the DNSSEC stuff, but the name server is also part of it.

PAT KANE: Like A records, MX records - a whole bunch of things could fit in this same model. It's very interesting.

JACQUES LATOUR: But just mostly a delegation, right? DNS.

EBERHARD LISSE: Okay, last question?

DAN YORK: Sure. Dan York with Internet Society. You know, Jacques, because I'm on the list, that I think this is a problem that has to be solved, and

whether this is the right solution I don't know, but it's certainly a solution, and I'm glad you and Paul and others are working on this, because this is key. We do need to automate this so we can be able to make it easier for DNS operators to work with the registries to make this work.

Two questions: one is what are you doing to socialize this out to the registrar, community beyond .ca? Have you been talking to other registrars or other TLDs?

JACQUES LATOUR: We're going to write something in Circle ID about this. Then we had the IETF Meeting, I think was EPP, a Registry EPP Session we had. There's more…

DAN YORK: Sure. I think as you get that out there, and now that you've got a demo out there, I would certainly be glad to help. Let's talk about how we can help promote this in some way. The other thing I'd mention - you put up a list at the end for an email list. I'd mention to folks there is a specific one which, if you just do a search on "DNSSEC-auto-ds" there's a mailing list specifically for people who want to work on this problem that's there.

JACQUES LATOUR: That's the list you want to use?

DAN YORK: Yes, that's one that's specifically focused around this, that people can subscribe to. They also can use this one too if they want to. Anyway, I'd encourage people who are in this, listening to it, to join with us and try to look at how to address this issue.

JACQUES LATOUR: There are more people on that list than the other one, so I'd rather use this one. I'm just saying.

EBERHARD LISSE: Okay. Thank you very much. Next one is Maurice Oviedo from .cr.

MAURICE OVIEDO: Hello everybody. My name is Maurice. I'm with the .cr team. First of all I wanted to thank you all, and thank you for the opportunity to share part of the work that we've been doing now for a while, with such a excellent group of experts, and hopefully we'll be able to get some input on what we're doing, improve it and make it better. That's basically what we want to share. There's a natural process that every department of every organization usually follows, where you constantly assess what you have, look for a way to make it better, you implement it and then yo do it again.

So in our case, we started to check what we have, define a couple of goals, what we wanted to do, and have been working on that for a couple of months. We wanted to basically continue adjusting our systems, so they are highly secure, to be fully tolerant, fully distributed

and to be economically feasible. We looked for the combinations for how we can get these points running and to continue improving what we have.

In that process, what we did was identify a couple of points we wanted to check. The first one was to analyze our existing infrastructure, and we wanted to know how we can use it better; what did we have to do to use it better? The first thing we did was basically identify some areas for improvement. The first one was to have a better leverage of existing devices, to take advantage of what we had. The first thing we managed to do was utilize the environment that we had.

We had to select a virtualization environment and run some tests on that, and we'll talk about that in a minute, and then we also had to adjust the existing services so they can use the new platform that we were going to deploy. At the same time we wanted it to be scalable to adapt a new project. One of those projects was, for example, to have a full size replication, so that in case of a major failure we can have not only the [curricular 01:19:15] assets on a different site, but have full operations of the NIC in case something happened.

Those were the key points we were trying to accomplish or the areas of improvement we wanted to cover. The first thing we did was to select the virtualization platform. We selected a very interesting one. It's called GANETI. Probably some of you, most of you, already know about it. It's a platform that for us is very interesting because of the features it has. Basically it's a cluster virtualization management system, which means that it works on the top of Xen or KVM. It allows

to facilitate the administration of the [brutal 01:20:07] servers and administration of the fixed physical nodes.

So it's quite interesting the way it works. It was designed by Google, and it was designed for Google. They wanted to use it on their office infrastructure for running their servers, and they made it open-source in 2007. So since it's open-source it's a [rating 01:20:35]. It's not necessary to invest in licenses. So if you have limited resources or a small budget, it's quite easy to get it up and running without making a big investment. On the other hand, something was very interesting on the way it provides the high availability environment.

Basically you can have different storage options within the cluster. You can have file based storage, plain disks, using LVM, or you can also use DRBD's synchronization within the hard drives. So how does it work? Basically you create two instances, or replicate an instance in two different physical servers, and you can get a right one, high available environment, through the network. You exchange or synchronize both disks on different servers to have high available environments during that work. So it's very interesting.

You can start with a single node and then scale up quite easily. You can go ahead and add different devices. They don't have to be the same specs. It's quite easy to grow the cluster and to be able to exchange information between those different nodes. Something that was also very useful is that you can have live operations within the cluster, which means that you can move virtual servers between

different physical nodes, and you can do it live. You have [serial 01:22:25] downtime.

You can plan maintenance for a complete node with a bunch of [unclear] machines, and you just need to migrate them between different nodes, and that will work without any downtime. Same thing with backups and live operations. So it was really good for us and has been working quite nicely. The way it works, this is basic deployment of GANETI. You have a cluster node. The node will be the physical server. You create instances, which are the [unclear] machines, and then you can replicate, or they are connected either through VLANs Trunks or Open vswitch, or any other platform.

So the instances by themselves can be replicated between the nodes using DRBD. At least that's the way we choose to do it. At the same time, this creates the GANETI cluster where you can manage everything. So that's a basic deployment. It's quite scalable. It can scale up to, for example, this SYNNEFO platform, which is this one that is here, where you can have different nodes, which are controlled by different clusters. Then you can mix the clusters and manage them through the cloud, and then you can create an API integrated to your systems and it scales very well.

The people from the Research and Network Institute, they have ten clusters with more than 6,000 VMs, and that's really cool. It was an optoin that we saw that was very useful and could be adapted to our needs, and at the same time had the opportunity to grow. This is the deployment we were able to come up with. It's a deployment where

we use the infrastructure we already have. We didn't have to buy infrastructure or make changes.

At the end of the day, the idea is using the exact same servers that we already had, just configured differently, at the same time something interesting is that if you want to grow it's quite easy because you can have centralized storage with a [SAN 01:25:08], but it's not necessary in this case. You just need to add a hard drive space and RAM in case you want to grow the devices. What we have is a production cluster.

In this cluster we have different nodes, and we are running all the services already on top of it. So we have the registry services, the office services, we run the DNSSEC, that we'll talk about in a minute, and we also run the Internet exchange point, so it's on top of the node as well, and a couple of other web services and so on. We also had the opportunity to create a lab cluster so we can test everything out, and our development cluster.

Since it's quite easy to have it running, it gave us the opportunity to replicate, to move information between the clusters, so that we can transition the new services or the new changes to different stages, and to make it more secure. Then we are in the process of installing an alternate site cluster. We're about to have it ready, and we hope to have everything up and running by the end of July. All the team is working hard. That allows the possibility of, in case of major failure, to be able to run the full NIC services in an alternate site.

So the next point was once we had the platform that we wanted to use, the transition of the services, some of the core services that we were

running were first of all the registry system. Whereas in this [unclear 01:27:08] people from .cc, the Czech Republic, FRED. It was deployed as a centralized environment at the beginning, and worked perfectly for a while. But with this new platform what we saw was the opportunity to decentralize the different components. We distribute the different components within different servers.

It allowed us to create different security policies, depending on the type of access. For example, it was not the same thing to access the WHOIS model than to access the database. So being totally separate from one another helped us to increase the security policies. Every different server is high availability as well, and different approaches are taken depending on the type of component. Then it allows load sharing in the way that the load is separated within the servers, given the main function or the main servers that they are provided.

We were able to do it without disruption or downtime. It was very interesting - during the migration, we were able to have everything up and running without any disruption, and it was quite easy to work with GANETI, in order to migrate from other platforms like BMWare, or any other one, to get it up and running. This is a simple screen of how everything is looking. We have the FRED registry system separated into different components. Each component has its own policies. We have a main registry interface, and [registry interface 01:29:09] as well.

It works through EPP connections with the main system. Once the information gets to the FRED system, basically we generate BIND zones. We move them to the DNSSEC signer. We sign the information

and from then it goes to the hidden master and it distributed to different secondary servers. The other servers that we were migrating, or was critical migration, was DNSSEC. .cr has been running DNSSEC since 2012, and it was very interesting how to use the new platform, how to provide high availability, and also how to use the system that can be shared with our own customers and can be cost effective.

So as with any other service, we continue to assess it, and the requirements are basically a system that's very secure, efficient, highly available, can be shared with the customers, well documented, to have the possibility to create backups in case of failure, so we can be up and running quickly, and it can be updatable. That's what we wanted to do. We came up with a system that's based on smart cards to create the keys, so we're starting to get up and running.

This process migration started with a workshop that was held in Costa Rica last year, in April 2014, where we got the help from ICANN. Richard Lamb was there with us, and also NSRC. They were the instructors. It was a very intense week that led to this that we already have. We started working with the tests and checking the different combinations to get up and running, and it was fully deployed by October 2014, and since then we've been using this mechanism. What we do is create the KSKs and ZSKs, to generate it through the smart card. We do it totally offline.

Once we have the keys generated we pre-generate signatures. We create the key bundles offline, material for a determined period, and then the bundles are exported to the signer. It's quite interesting,

because the key signing key network never touches the net. It's totally offline. Once everything is signed in that ceremony process, the KSK is kept safe and never touches the net. We're using 2048b keys, basically on the KSK and ZSK. We're using them for both.

We're using this modified version of the tools and scripts that Richard Lamb created to work with the smart cards. It's quite interesting, because we adjusted it to our infrastructure, this information, and also are using DNSSEC signed zone in order to make the sign on the server. It's basically a bunch of scripts that were modified, and they're working very well. A couple of details that were part of this project are we were able to create full key signing ceremonies. We're doing two different signing ceremonies. One of them is for .cr, for the TLD, and then for the subzones we have a different one.

Both of them are with the integration with the local community, so we have a full ceremony on that, and that's part of the trust that we'd like to generate with the process. Then it's a very efficient process. Right now, to sign everything we have, it takes around 20 seconds, and the interesting thing is that since it's on the server the load can be shared between different cores within the server. So if you have [virtual 01:34:06] cores it will be faster. With one core it was taking around a minute, ten cores taking 20 seconds, so it depends on the capabilities that the server has. Again, you can split it depending on what you have, and can be scalable.

The smart card signing by itself is not fast, however since we're just generating the keys and signing some keys with it, it's just done

through the key generation ceremony, and then the zone signing is done on server. As I mentioned, the KSK never touches the net, so it's very secure as well. This is basically how it looks. After generating the KSKs we export the key bundles up to the DNS signing server, which is isolated. It's a one direction process to move the signed zones to the hidden master, and then from there we go to distribute authoritatives as well.

Another thing we've been doing to promote the stability of the system is continue working on the distributed DNS for the secondary servers. I want to just quote Martin: "You never have enough Anycast," so it's important to continue building the stability using different clouds. In our case, we added PCH to our network, and that worked excellently. We had the opportunity to get presence in every continent and major IXPs around the world with this move. Then we're also using ISC and RIPE Anycast clouds.

We also have physical servers in Costa Rica. Our friends from NIC Chile and NIC Mexico also serve as secondaries for us. Estimated around 70 servers deployed around the world what we're using, and it's provided a lot of stability as well. Our [unclear 01:36:34] we are participating - we're very excited about it - it's working with LACTLD on the project of the Anycast network that they're creating. We're participating in two different approaches. One of these is by being a user, to use the services to replicate, and then the other one is to be a node; to host a node from the Anycast servers.

This leads us to the last point on this slide, which is that we also run Costa Rica's national IXP. We have of course .cr connected directly to it, but working with LAC TLD allowed us to connect that node, and will allow us to connect directly to the different IXPs, which are already part of the IXP project. It's called CRIX, in case you wanted to check it - CRIX.cr. We started with it in 2014, and already have 16 different members, 17 with us, and it's growing quite fast. It's very interesting to make an [parentheses 01:37:53] here, to have both sides of the table, to be sitting on both sides.

Being working at .cr, as the operator of the .cr DNS, it's also interesting to see how it looks from the other side, from the ISP side, to get to know better the community and to be integrated with them. so it's very interesting now that we're also managing the IXP to have all the ISPs sitting at the same table, and to run different initiatives. Some of them are related, for example, to DNSSEC. To be able to work with them, we're creating a Workshop that will be out soon, hopefully, which is specifically designed to teach DNSSEC to the ISPs and to enable validation on their recursives. It's very interesting to have them on board and to be able to have a more direct channel within Costa Rica.

This is basically how it looks. Hopefully we'll be able to add LACTLD soon to the distribution system. It's not that updated, but it looks something like this, the distribution within the Anycast cloud - different clouds and servers that we're already using. Just putting this out together, this is something that puts it in a different perspective; to

work within the security, stability and resiliency of .cr, Different approaches are being taken.

We are working with the distributions, with [unclear 01:39:52] platforms, security, with the different DNSSEC models, also with our new website, it's adding security to registry systems as well, and with the evangelization of the different things that we're running, DNSSEC, IPv6, DNS itself, we're doing a couple of things - workshops within Costa Rica, participating in national events, talks with universities, colleges, international events like this one, and advisories. It's part of the [unclear 01:40:28] we're trying to create, and I just wanted to put a different perspective on how NIC.cr is working, and what it's doing right now.

A couple of conclusions - the improvement cycle never stops. You continue to assess what you have, to improve it, and to get the input from people and experts like you, and to start again - to implement the new things and continue growing. Improvement doesn't mean that you need to invest a lot of money, and that's something that we wanted to mention here. In our case, gaining a high availability system didn't mean investing a lot of money. It's basically getting the correct combination of the tools that are probably already available, and at the same time to keep asking the proper question, asking for ideas from the experts and people there who also different networks.

Important to prepare the network to grow, and getting the input is one of the best things that help us come up with this new model that we're

using, and I hope you find it interesting, and your input is highly valuable. That's basically what I had prepared.

EBERHARD LISSE: Thank you very much. It was a really comprehensive overview about how to run the registry. Any questions? Okay. Thank you very much. [applause] Almost our last speaker is Adiel Akplogan. Is he in the room? Adiel has been appointed, a few months ago, as the VP for Technical Engagement. He is to be the CEO of AFRINIC, the regional IP Address Registry, and I've asked him to come and say a few words and see what he can do. As you know, my hobby horse is outreach to smaller, especially developing country TLDs.

ADIEL AKPLOGAN: Thank you Eberhard for giving me this opportunity. Sure, when my position was announced, most of the questions I've heard from people is what exact technical engagement means. ICANN, being a technical organization, why is that important? It took some time, even myself, to wrap my head around that. Of course, my presentation will be a bit different from what you've had here so far. It will be a little softer. But I think this audience is the best I've seen within the ICANN Meeting to expose, exchange and share some of my idea about technical engagement and how that could be very important for ICANN.

I will take you through a few slides. I'm not going to spend time on all of them, but we'll focus on a few key areas of this slide. What I will expect as well from this presentation is some exchange and input from

you on what you expect from technical engagement from ICANN, and how we can work together to make this very key area of ICANN's role and responsibility what it should be.

I will talk briefly about the vision and the goal of technical engagement as we are defining it. It is for us a tool to develop, maintain and sustain ICANN engagement with the global technical community. This is specifically in line with its own strategy objectives. When I started, one of the things I've tried to do is see, as a technical organization, what exists with ICANN. What I've noticed is that ICANN engage in a different way from different corners of the organization, and engagement happens broadly within the organization.

But probably when you're outside - and I'm using my own experience as an RIR - you don't see that aspect. What we see more often is the Internet governance aspect, is the policy development aspect, and even on those aspects, that goes a lot towards the domain name business, while we know that in the ICANN name there is two N's. How do we make sure that ICANN's strategy that's there and include everything?

The other aspect of this technical engagement is how do we make sure that ICANN constituencies are constantly and continuously exposed and aware of what is happening in the technical ecosystem globally. That is a two-way engagement - one from ICANN to the external world, but also from the external technical community to ICANN. I was just sharing a session right now where different SOs but also the IETF in the area have presented what they're doing. Most of the activity doesn't

happen within ICANN or during ICANN Meetings, or within the ICANN framework.

So how do we make sure that this community understands and is aware of what is happening in this area and participates? Roughly this is the goal that we set ourselves, and as you have seen from Fadi's presentation this morning, technical engagement would be done now through a newly set up executive team leaded by David Conrad, who is the CTO of ICANN. I would be supporting him there from the Global Stakeholder Engagement, but mainly working on strengthening this relationship with the other parts of ICANN.

A few prerequisites we have identified as key elements for this engagement, first of all, is to have a clear technical engagement. As I mentioned, we have identified more than 14 points of our engagement within the organization. How can we make that more coherent so that we improve its visibility beyond just the technical service that ICANN provides? How do we clearly make ICANN technical activities and technical engagement as one pillar of ICANN's role and ICANN legitimacy?

The second element is better coordination, which is a follow up of that, of our technical engagement. We should do this in an orderly manner. That is also something that we are working on. Provide the community a single and consolidated point of entry to the technical space. This will be translated into many areas.

The first one is the creation of this executive team made up of technical experts within ICANN that can be continuously consulted into

overall ICANN decision-making processes, which is already an important step, but also in ICANN community or outreach; having a very well-defined area where if you are from the technical community you can go to and have the information you want.

For instance, on the ICANN website it's not always easy, if you are from the technical community, to find information that is related to us. We will be working to make this more effective. Our contribution as an organization, but also as a community, to different technical forum like the NOGs, regional, local NOGs, the IETF and others, we attend, we go to those meetings, but as ICANN, an organization, how do we effectively contribute our experience to the work of those different technical forum?

Open technology, open-source, this is something very key for an organization that is at the head of the Internet protocol and security and stability. How can we better promote open technology and best practices as a technical organization, and finally strengthen our partnership with the I* organization, who are an organization that don't come always to ICANN Meetings, but their role and activities have a huge impact on what we do here.

As I mentioned, we have identified 14 points of engagement areas within the organization, and in the ecosystem. Some of them are related to policy development and their implementation by ICANN. You will see the other SO. You will see there some internal function of ICANN itself, like its Information and Innovation Office, the Security and Stability Office, the CTO Office, et cetera. Those activities are very

focused on providing service to support policy development related to this activity, but also implement them.

Beside that, we also have other areas of technical engagement, like the one with the I* plus - I call it that because I add the Network Operator Group - and other Internet governance related organizations that are related to technical areas. I listed the ITU-T/D, the IGF and Best Practice Forum, the OECD, et cetera, where ICANN should get more engagement by developing cooperation and partnership with some of those organizations. But those partnerships shouldn't be just a theoretical partnership, but based on concrete projects and activities that can be developed together for the good of the community.

Finally, we have what I call the non-conventional stakeholders, that are an organization that I'm not familiar with what ICANN does, such as a professional association or industrial association. Sometimes we can extend it to other organizations like IEEE, GSMA, et cetera, where the awareness or engagement activities more about awareness let them understand what ICANN does, which is a bit different from what I can say about the I*.

One of the things that I will talk about here in relation to our engagement plan for this specific audience is the Tech Day. It's one of the key elements of our engagement plan. I think it is the only formal forum during the ICANN Meeting where the technical community get together and exchange on issues that are of interest to us. One of the thinkings we want to share with you and get your feedback on is how do we make the Tech Day something broader.

How do we make it something that has a footprint that's wider than what it has today, evolving it from the ccNSO Tech Day - I think that evolution has already happened in some extent - to an ICANN or community Tech Days, with an S at the end, whereby we use the opportunity of the ICANN Meeting to exchange more, not only within the technical community but also from ICANN to the technical community on several activities that happen internally within ICANN, where we can share our challenges, where we can get feedback from the community on those activities, and where we can also develop a stronger partnership in that.

The Tech Day is an area where I'm sure we can do a lot more to position ICANN technical activities and also strengthen the legitimacy of ICANN in the technical area. Capacity building and partnerships with other organizations and capacity building is another area. Also making sure that the SOs that are related to technical engagement, like the ASO, and ACs like the SSAC, RSSAC, TG or ISPCP, are very well know by the community outside of this room, to be able to participate more and to get their voice heard. What I would like to hear at the end of this presentation is how we can do that with your support, and what other areas you think we are missing or can put more emphasis on.

I think these slides will be available for you to look at later on. I'll also touch on this. For the past three months or so we have been focusing more on the internal aspect, understanding what the machinery is internally when it comes to technical activities, technical engagement for ICANN. We are now in the place where we are reaching out to the community widely to get feedback and to start implementing some of

the findings, some of the suggestions we are making. As you have seen, some of them are being materialized already with the announcement made by Fadi today.

We will continue this. We will continue working as well with the communication team internally to make sure that the technical aspect of ICANN activities is always represented and highlighted in our communication strategy. Another very important aspect of the technical engagement is the decentralization or regionalization of this engagement. I have the privilege to be working within the GSE Team where we have several Regional VPs that are heavily engaged on the ground, at the regional level.

We believe that if we want to be effective in our technical engagement strategy as well, we have to rely on the regional technical community as well, using different opportunities that already exist, and try to build partnerships around them, such as the RIR, who already, by the way they operate, have a very regional operational model. We have ISOC Chapters around the world. We have NOGs, regional and local NOGs, and ccTLD regional organizations that already engage in some aspects technically at the regional level.

How can we as ICANN build up on that to be more present and effective in our technical engagement? That goes down to supporting ccTLDs regionally, ensuring that their relationship with ICANN is up to the expectation; both technically but also on the policy front. That is the essence of these slides. I'm not going to spend time on all the

slides. I just highlighted the areas I think are relevant to this community.

To us, technical engagement is something that goes beyond just the engagement but also involves some kind of mindset change, not only within the organization, but also within the community that's very familiar with ICANN on the fact that ICANN is one element of the overall technical ecosystem that sustains and maintains the stability of the Internet. We have continuously, when we engage, keep that in mind, and keep in mind the original role of ICANN as the coordinator of the different elements of this system. Thank you.

EBERHARD LISSE: Thank you very much. Any questions?

JAY DALEY: Hi. Some feedback for you Adiel. Thank you very much for that. You talked about Tech Day becoming a place that it can improve the legitimacy of ICANN's technical function and things. I think it is important to note that Tech Day currently, nobody has any priority to come here and speak. All the talks are judged on merit, and I think we need to ensure that ICANN doesn't feel that it would have any priority here to come to speak. That's fundamental. If it does become then ICANN-specific, that would be a problem to us.

I think that in terms of dealing with the I*, we need to avoid doing that, because there are plenty of people here who go to all of those things, and if we go to all of those things and then hear about all of those

things at all of those things, then we're never going to do anything except hear about all of those things. We need to avoid that. Also, we have traditionally succeeded quite well by individuals who have done a piece of work in another area, coming here and bringing it to us. We've never really had that much involvement in many of the RIR policy stuff, largely because it is so phenomenally boring, none of us want to have anything to do with it really.

We're all part of the same family, but we don't want to overdo that bringing of things together. Finally something we talked about the other day as well, we have a very strong technical set of presentations that people bring, based on their experience, and we also need to be very careful about moving that any way into an education role. I think there needs to be an education role, definitely, but that can take place elsewhere, simultaneously with this, because we want to maintain the technical ethos here.

ADIEL AKPLOGAN: Thank you very much. Well noted. I will comment a little bit on that, but I think that's a very good comment.

WARREN KUMARI: Warren Kumari, Google. I think one of the things we need to be very careful is that ICANN does outreach by proving competence and not just showing up at Working Groups or other organizations, I* things, and saying, "We're from ICANN, we're here to help." Actually demonstrating that ICANN has capabilities and can contribute in a

useful way I think would be an important first step. Currently many groups don't have hugely great respect from ICANN, and I think there's going to be a fairly large culture shock with many cases.

ADIEL AKPLOGAN: Let me start by that caution - that part of your comment. I fairly agree with you. I'm coming myself from the area of industry, and I know how I used to see ICANN from outside. I understand that. I think that is the essence and the bottom line. I mentioned somewhere in my slide that we want to participate in those things, but not just participate - we want to contribute. We want to be there as part of the community and share things that we do. Being able to get that feedback and improve them openly, that is part of it.

I know it won't be an easy task, but let me share with you my own experience. I've met some very well knowledgeable people at ICANN about what we are talking about, and who really wants to do things. One of the things that was probably missing, and which is the step that we start seeing, is the proper alignment and coordination of this with the overall strategic objective and goal of ICANN in this specific area.

WARREN KUMARI: I think that ICANN is making some good steps in that direction, with people like Terry Manderson, who's now an IETF AD, and some of the recent hires as well, I think will help some.

EBERHARD LISSE: My own view on this is that we have seen very good presentations from ICANN staff today. We have had Ed Lewis and the L root people, Kim Davies was not really technical but it was good that he came, and we asked for this again because we are all affected by this. But on other occasions we have asked ICANN to present about issues that we noticed, for example when the website was hacked, and they just didn't respond.

In the end , the focus of this has always been to try to share experiences; what works and what doesn't work, like this morning, with especially smaller and medium-sized ccTLDs, and now also gTLDs, and that's what my two previous speakers said - we should carry on doing this so we actually achieve something.

ADIEL AKPLOGAN: Definitely, I couldn't agree more with you all on this. It is a journey that we've just started, and we will be working with you to ensure this more and more. I think by already having a stronger technical presence at the higher level of the organization is one step already for an organization like us, and now we need to strengthen that, to make sure that we sustain it, and we also go beyond.

I think the Tech Day is one step, but going beyond that as well so that broadly when we say, "ICANN Technical Team," we know that they can all speak with confidence and also engage as much as they can, and they need to ask all of us as well. I'm also putting a lot of emphasis on developing projects and different activities in cooperation with different organizations. That is also a way of improving that

engagement. You will be hearing from myself, from David and from all the Members of the Team on this. I got the other comment about not overdoing this, just by speaking [unclear 02:06:50], but by actually having on the ground some critical and key activities that support this.

EBERHARD LISSE: Okay. Thank you very much. [applause] Now we come to the really boring part of the afternoon, mainly the PGP signing ceremony. How many of the participants have PGP or GPG key? How many use it regularly? Less than that. How many of the regular users have their key signed by more than one other person? Very good. For the other half that haven't done it, we all know what PGP or GPG keys are. It's cryptographic material, private-public key, to which is [appended 02:07:50] user identification packet, as they say.

It works usually that your email software automatically signs your email if you configure it, and incoming email is signed or encrypted, or automatically verified or decrypted. The point however is how do we know that the person having access to my key is actually me? If I feel my key has been compromised I can remove it, but the idea is that Warren Kumari doesn't really know that el@lisse.na is actually Eberhard Lisse at .na. That only works if my key has been signed preferably by somebody that Warren knows, or that Warren's key has been used to sign.

In other words, the more keys get signed, eventually a web of trust occurs where if one key has been signed by one of the individuals in that, subsequent keys are at least identified as being persons, and if

the same diligence is provided that each individual user is to sign a key, then one can say if Warren has signed Andre's key and I signed Warren's key, then I trust that Andre's emails are from Andre in fact, when they are signed as that.

Now, what we have done is we've created and advised in our emails repeatedly on a website called Biglumber a key ring. Robert Guerra, who is probably not present, he's on NomCom - so he's much worse off than we are as far as meetings go - he has set up this key ring, where we can load all our keys. We learned one thing, that the criticals, like "check error" would be a critical, is not really accepted very nicely. That's not generally speaking a problem, because you can tell GPG to use UTF-8, but your key ring didn't do it so we had to work our way around it.

We have 20 keys uploaded and only about ten are here. [unclear 02:10:19] is here. There you are. For each participant we have generated a worksheet, and for each participant I have generated a single piece of paper, which usually every individual has to identify him or herself to every other participant by showing open ID. Because this is a bit of a nuisance, what we have done is I've written a little script and generated a worksheet where every participant can tick off whether he has identified the other person and whether the key is correct.

The way this is going to work is we take the key that we have generated, put the ID card, and put the scanner on top, so everyone can see it, take it off, we read the key. Everybody who puts his key on

the table must quickly check from a separate source that the key is correct, and then when you compare that key to the worksheet that should be okay. In order for this to work, I need to have a little preparation here. I've brought my little toy. We accept official government ID cards. Everybody is free to accept whatever ID card he wants to accept.

Generally speaking, the tradition is we use a government-issued ID. As a reference to good practice, as we experienced in Singapore, I had a key that was that old that it broke a lot of software. [background chatter] Can everybody see my key? Can everyone read it properly? I will read it to you. Okay. Can everybody see my ID card? It's a government-issued ID card from my own country. Participants need to take their worksheet. On the left side of my key is number six, so on the worksheet my key is number six. I've checked the long key written there from a separate source is correct.

I am now going to read my fingerprint. This long key written on the bottom is supposed to be checked against a separate source, and I myself have checked this morning it's the correct key. Make sure on your worksheet that this identical to what my key is written: 7399 BE0B AEFB 4AE5 EDB9 4A54 9705 1DA3 7945 3FAB. If that corresponds to what you have on your worksheet, you check on the worksheet "fingerprint checked". If you have seen my ID that that ugly mug shot is my face, you can also check "identity checked".

When this is all done, you go through your list, then download the key you want to sign, sign it, upload it to the nearest key server and then

we'll propagate. Who wants to be next? Warren or Jay, doesn't matter to me. The number is listed on the page, so it's easy to see.

WARREN KUMARI: Hey everyone. That's me. My key fingerprint is E712 739B 73F3 7206 39DF 8124 75C2 1930 E4EE 98DA. If anybody doesn't like my driver's license, I'll have my passport with me later.

JAY DALEY: Jay Daley, number 14. That's my passport. It's B5E5 1DC1 5651 710C 22D9 B716 202C 6910 74D1 5853.

WARREN KUMARI: Another best practice is while you're checking these off on these sheets, you might also want to sign the bottom, somewhere on the sheet, so that way you know it's your piece of paper and not somebody else's, who hasn't actually checked or decided to swap in a different one.

EBERHARD LISSE: Which one are we going to take?

JOE ABLEY: There are two keys on here. one of them on Biglumber, I can't seem to update it, it's expired, but the key servers have the correct one, same ID, so I don't know. I'll read them both out. The first one is number four: 80B7 8D10 922C ED01 CBE9 85D8 348F 0CBD 8652 3A2C. Both of

these I checked against the laptop. The other one is the Dyn key, number ten: A80C EC8A A35C DC5C 08F5 F3B1 D721 1DAB 8668 7134.

SPEAKER: This is number 16: ACBA 1BA3 28D4 EA26 42A5 3244 3B41 8E9F D88A A9883.

EBERHARD LISSE: If anybody cannot read the numbers or has issues with identification, raise their hands so we conform this best practice. That's a Namibian ID card, but I don't know why it…

BEN FULLER: Okay, I'm Ben Fuller, number 17: E9C6 3F9C 1076 3FA3 1B08 EAD2 1CDA A593 F319 442A.

SPEAKER: The number is 19: 697D F467 885F B1D0 EB73 EB51 812F C941 CEDA D830.

EBERHARD LISSE: Who else is here? [John Kane 02:25:35] was supposed to be here. Joao Damas can come up. Patrick Faltstrom is not here, and John Levine.

JOAO DAMAS: I'm Joao Damas. That's my ID. The fingerprint for my key, which I actually checked against the one in my laptop, is: E45D 489A 334F 71BE 4DAE BCE8 95F5 5639 FFF9 9715. It's number three.

SPEAKER: I'm [Ross Algado 02:27:15],number eight. My fingerprint is: 355A 539D 1488 EF91 73DC 56C6 D770 BB46 F120 A230.

EBERHARD LISSE: What is the best practice requirement with regards to identification? Is a photocopy acceptable? What we are saying is we read the key and whoever accepts the identification ticks it as identified. There you go.

WES HARDAKER: I'm number 21. The local advice for carrying your passport in here is you don't, you put it in your safe, so that's what I was following. So I have a copy of it. That's my copy of my passport. I will read the fingerprint. Sign it if you want. If you don't want to, because it's a copy and not my real one, that's fine: 6248 FBFF 6300 5E0F E5D7 776C 3040 7DC2 E578 7FD6.

EBERHARD LISSE: No-no. It's not what I do. It's recommended best practice.

MAURICE OVIEDO: Mine is number 15, Maurice Oviedo: 4718 295A 38BA EF55 1110 AC75 F438 3165 577F 134E.

JORDI IPARRAGUIRRE: Jordi Iparraguirre, number 18: 4759 C4F8 FE7C BCD4 3E4F A926 7B3A 39C1 C413 D990.

EBERHARD LISSE: It's Andre of course, no?

[ANDRE PHILIP]: I'm [Andre Philip], here's my passport, and just a warning - Mikey was not able to [unclear] Biglumber, so if you want to sign it, please don't [unclear] from the key servers. My fingerprint is: 6D7D 0C11 6BB5 A207 CFEB 5C9E 3F36 0F84 480E 858F.

EBERHARD LISSE: As I said earlier, we had issues loading this key up to Biglumber, because Biglumber doesn't accept UTF. His key you must load… If you have Mac, it's very simple. You go to the PGP key ring. Just type the short number on the left side in, which in his case would be… If you look for that key on the key server, offer the email address, the key will come down. They you just check it has the same fingerprint and then you can sign it and upload it again. The plan now is that everybody who's had his identity verified sends every other key of each individual for whom he or she has verified the fingerprint and identification.

If somebody is unhappy with the paper copy of identification, he doesn't sign Wes Hardaker, I'm perfectly happy with that. The protocol

is that you must be satisfied with the identification, so it depends on each person. The idea is now on Tuesday, Wednesday, Thursday over lunch, I will go to [unclear 02:34:32], and I would appreciate if a few of us still come there, because if one or two stragglers come then we have enough people to sign those keys. Otherwise it's just me who signs it, and that's a little counterproductive.

If anyone in here wants to sign the key, please download the key to Biglumber and email me, because I have to then reprint the whole worksheet and things like this. Thank you very much. Usually I have one unsuspecting suspect give us parting remarks, but I think since we've done a practical at the end, we can dispense with this for the time being. Thank you very much for coming.

[END OF TRANSCRIPTION]