Eberhard Lisse: Most of you will know me, I am Eberhard Lisse. I am the Chair of the Technical Working Group. We do the usual Tech Day on Monday's, even later in the morning when it comes close to lunchtime because it's sponsored. This time there will be a (inaudible) block here. Today we have decided to focus a little bit on IDN and on IPv6 so we have a few DNSSEC topics, but DNSSEC has been more or less beat to death. So I'm not really trying to focus too much on it.

In the afternoon I want to talk a little bit about a business continuity situation, in particular about communications after a natural disaster or after a big disaster; and then we'll have a register ASP give us an idea about some studies about scaling registries they have done. And then we'll have a question time roundtable that is without any set agendas. We'll have a few interesting people sitting here. I'll have Alejandro Pisanty, Jothan Frakes, and who else? Two others; Nigel Roberts is going to manage any questions you may have on this. Let's hope we get a bit of a discussion going.

At 4:00 we have to sort of relinquish control, but we won't leave the premises because the DNSSEC for Everybody session is starting here. Whoever wants to participate is more than welcome to remain. I encourage everybody who is not very familiar, who hasn't done it, or who thinks he can benefit to do that. It's not run under our umbrella; it's just the same room, so I put it in different font at the end just to remind everybody that it's happening.

Okay now, as you may have read ICANN has finally filled the post of Security Director with a fairly eminent individual and I'm very pleased to have been able to get Jeffery Moss in to give us a few thoughts. Invited speakers, as I say I don't usually give a topic; I tell them they can talk about whatever they would like to and I'm quite sure he's got some interesting remarks to say.

Jeffery Moss: Let's see if this works. Hello, good morning everyone. I want to thank Eberhard and the organizers of Tech Day for having me here. And you get the dubious honor of this is the very first official public participation of myself in any ICANN meeting. So you'll be the beta test guinea pigs for what I have to say.

I come from a technology background from consulting and starting tech businesses and having tech businesses fail. So I come from a small – medium business background. And then over the course of my career I find myself being drawn more toward policy and organization.

So my perspectives have changed. I used to always believe there was a technical to every given technical problem and I'm starting to realize that policy plays a role there, more and more than what I had acknowledged previously. And so the attraction to ICANN as a giant policy coordinator was more and more interesting. So that's one of the reasons why I took the role.

In looking at ICANN from the outside, and then hearing stories from people like ccNSO groups who they don't have a contractual

arrangement with ICANN, but they participate, it's very encouraging to me because I think there are very few things that ICANN does in security operationally. The DNS OPS Groups operate its L Root, they manage some DNSSEC zone signing and IANA handles some root zone changes. But I can't think of too many other things that they do operationally.

So I think primarily my view of the security function is a coordination and a collaboration role. And you can look at, for example, the collaboration with the Conficker Working Group or how sometimes we provide assistance and coordination to law enforcement who are just trying to figure out who they should communicate with in an investigation, where the valuable WHOIS data rests, which registrars to contact and so forth.

And in the collaboration where it's really about us trying to connect with other partners and getting partners to work with each other and in that area I would say we've provided to some of the ccTLDs, we try to provide training, capabilities training around DNS and I don't see that stopping. I'm actually trying to accelerate additional training.

I think the more training we can get out, whether it's to the cc's or to a private group, anybody that can better understand DNS and DNSSEC autonomous systems and the other ways in which the unique identifiers are used, that's good for all of us. So I'm going to try to continue developing training programs around that.

We collaborate with other groups around, not just capabilities, but continuity planning, disaster recovery, we participate in, in the United States we participated in a cyber storm exercise – I think some of the cc's also participated in cyber storm. We act as sometimes subject matter experts. We're consulted and participate in the IDN implementation on Variants.

So I see security really as an enabler for, not just you, but for the community at large where we want to help out and act as subject mater experts when necessary; sort of a way to connect the dots between different groups when necessary to help everyone operate more smoothly to enhance the security and stability of the name system of the unique identifier system. And it's really not to try to take over anybody's operations; there's enough other problems in the world for us to work on.

So I don't have a lot of prepared remarks to say, but I do want to extend an open hand and an open door and that I want everybody to feel comfortable contacting security, contacting myself with any concerns. And from my side, this being my first ICANN group, I'm just really interested in learning from you all. I've heard really good feedback that some of the ccTLD operators have superior best practices and it's almost like an ecosystem where many different approaches have been tried and I'm interested in learning about many of these different approaches.

Honestly I'll probably end up borrowing some of them and championing some of them. I come from a world where the best idea wins and I'm not particularly interested in the politics, I'm

more interested in what works. So with that said, I don't have too many ICANN anecdotes, I just have this basic introduction. I am perfectly happy to take any questions and answers, give you any insight that I may have, but this is truly the beginning. I think I've been at ICANN about 30 days.

Eberhard Lisse: Okay, thank you very much. Do we have any questions? Come on, we can't let him off that easy.

Jeffery Moss: Yeah, you can't let me off that easy. I'll sit up here until I get some questions. Yes.

Eberhard Lisse: Please even with remarks use the microphone otherwise our remote audience won't get the benefit of them.

Jeffery Moss: But you have to figure out how to turn it on first.

Eberhard Lisse: The staff over there is very well aware of what's going on and they're turning it on for us. There is a question! Excellent! No question.

Jeffery Moss: I've got a ringer in the audience over here, Patrick Jones…

Stephen Deerhake: Steven Deerhake from the AS Domain Registry. Do you envision ICANN, at some point down the road, offering a type of audit service for if a cc is interested in say okay we think we're doing things correctly but can you guys come and look because you run L Root yourself and you have some familiarity with this?

Jeffery Moss: I see developing a sort of a best practices guide book so you can sort of self assess and self audit. That's purely a budget question I guess for resources, but developing a course where developing a self assessment methodology I think would be useful for everybody. So that's probably where I'd go first and then second if I have the manpower…How do you do that now? Does everybody develop their own audit or compliance guidelines?

Eberhard Lisse: Best practice is a word you don't want to use near a cc.

Jeffery Moss: Because it's "best" in each country?

Eberhard Lisse: No, because we don't have any obligation to follow any rule and then some of us are very territorial in that respect.

Jeffery Moss: But you must have your own best practices as defined by your…

Eberhard Lisse: As defined by whom?

Jeffery Moss: By yourself, by your government…

Eberhard Lisse: Yeah, that's the point. We have the larger ccTLD have got obvious are on the same level as common and that's not a problem. Where it is a problem is especially for the smaller ones and not so much for the middle, but the smaller ones have problems – they have down times, we can't reach them.

My neighbors in Mongolia are supposed to live somewhere in Portugal, try to register a domain and it's very difficult. I've got many patients in Mongolia, I see them in my practice, they give me their email to correspond and it's always in Brazilian or Portuguese and reaching then is just plain difficult.

We don't know where the guy is. We don't know how to reach it. I hear from other registrars that talk to us that they find it very difficult to reach some registries at all. So this is a very wide spectrum. What the plan was and is for this group is also to write down what works so we stay in contact on the particular things so we can have especially the smaller ccTLDs in developing

countries, and they have absolutely no idea how bad it is in some places as far as business is concerned. They just have nothing. To come here and to say best practices is not going to work.

Jeffery Moss: No. But they should have access to information.

Eberhard Lisse: I'm not criticizing you. I appreciate and I thank you for that, but the point is your problem with the cc's is that there is a wide spectrum and on the low end it's really, really bad. You cannot imagine how bad it is.

Jeffery Moss: I'll take your word for it.

Eberhard Lisse: It's very difficult. They don't come here. They cannot afford it, whatever. Or if they can afford it they go and do shopping, I don't know. It's very difficult. The information is there but it's very difficult to implement it. It's extremely difficult. Many of the ccTLDs in the developing countries are run by governments or by government institutions. Governments in developing countries is not as good as in my country most of the time; very difficult. And I come from a developing country.

It's very difficult to tell the government how to run things. South of Sudan is going to become independent and I've got quite strong

feelings for their situation; they're going to apply for a ccTLD and it's going to be very difficult. I've been there. Physical infrastructure but we have no idea how the management is going to be – who is going to get it within government, what's the policy struggle…

Jeffery Moss: But if a ccTLD does reach out and want advice we'd be happy to give it.

Eberhard Lisse: Sure. That's what I'm saying. The point is don't put your hopes too much that you can sort of line us up and line the cc's up and say this is how you do it and then everybody will do it.

Jeffery Moss: Oh no, no. I've been warned of that.

Eberhard Lisse: Don't get me wrong, I'm not saying ICANN isn't important policy, I'm just saying the ideal way of doing it you will find they say yes, yes, yes. And nothing changes.

Jeffery Moss: Okay, good. A second question.

Simon McCalla: Morning Jeff. A quick question - we've had a little talk over the last kind of – sorry Simon McCalla from Nominet. We did a little talk over the last year or so about the relationship between ICANN being operational and particularly around the area of certs and responders and so forth. What's your thoughts coming into the community on the relationship and how operational ICANN ought to be and how much it should work with the community?

Jeffery Moss: Well, I view it as primarily a collaborative role. So, where ICANN clearly has the authorization to be operational they should be operational and they should do that in a very methodical and stable way. As far as Certs go, I think some of the community efforts around there, I think there is – correct me if I'm wrong – I think there is almost is a community effort right now to sort of develop their own DNS Certs and I don't think ICANN really has much to do with that at all. And that's great. If the community comes up with something that helps us all then I'll provide support and I'll provide advice, but it's unclear to me that that's an operational role that ICANN should be pursuing.

Eberhard Lisse: That feeling is shared by many. Warren, you've got the microphone. Have you got anything? Anyone over there? Alright if there's no more questions, then I can…

Jeffery Moss: Move on. I'll be here all through Saturday or Sunday and if you have any questions you just want to talk privately, pull me aside and I'm happy to talk to anyone. Thank you.

Eberhard Lisse: Thank you very much for coming. Short and painless I hope. Alright, shortly before lunch there will be more I can tell you this because it's sponsored.

Next is Mr. Al-Fayez from Saudi Arabia – I almost said South Africa. He's going to speak about their IDN experience. I think Arab speaking countries are ideal to talk about this because they have a totally different script. So I'm quite pleased that he made this offer.

Now we first have the usual technical difficulties. Let's see if the techies get it to work. Not yet. It was working just now. While we are waiting, I'm expecting Ryan Tan today. Thank you very much. I didn't know, I wanted to make sure that you are there. There you go. You've got the floor.

Raed Al-Fayez: Hello. My name is Raed Al-Fayez. I'm from SaudiNIC in Saudi Arabia. First of all, I would like to thank ccNSO for giving me this opportunity to share our experience when launching IDN ccTLD .saudi in Arabic and I will say it as .saudi. My agenda for my presentation – I will speak about SaudiNIC, our experience and some statistics, what we have done exactly in the couple of years

ago and what is still coming as a project in the coming years. And the second thing I will talk about our experience in Saudi, what we have done and what is next. And then at the end I will share the lessons that we have learned when launching our IDN TLD.

So Saudi Network Information Center is a nonprofit organization operated by the Communication and Information Technology Commission – CITC. It's similar to IDA here in Singapore. SaudiNIC runs the .sa since 1995, the ASCII label. And the IDN and the Arabic label since May last year. We are leading the local community efforts towards supporting Arabic language in DNS. And we chaired both steering and technical committees for the Arabic Domain Name Pilot Project.

These are some of our latest achievements. We fully have IPv6 support as of January last year. And we opened registration for Arabic domain names under .saudi and this is an example, (inaudible), you can see it here. On May last year also we opened registration for domain name directly under .sa; so second level domain names was opened earlier this year, January 2011. And we updated our domain name regulation and procedures for submitting objection to Saudi Arabia, which is similar to the dispute resolution, in the local laws in Saudi Arabia, in April 2011, this year.

Coming projects: registry-registrar model. Hopefully we will open the registrars. And we hopefully will have DNSSEC soon. We are experimenting with the IDN emails because we know there are some issues, the standards haven't been yet finalized. Also we

keep testing IDN implementation and well known applications such as Firefox, Explorer, Chrome, etc.

We have around 25,000 domains registered and most of it under .com.sa; second .com.sa we opened this year. And for the IDN .saudi it's the third also. And the rest of the sub level going on. So what we have done for .saudia – when ICANN opened the Fast Track, on 16 November 2009, we immediately applied for our label and ICANN approved the string in 20th of January of last year.

And on the 24th of January we applied for IANA delegation and we have got the approval from IANA on 22nd of April. On the 5th of May IANA added .saudi and that you can see in this slide, to their own servers and on that date .saudi was operational. We have immediately tested domains that were enabled the moment that IANA added our label to their zone file.

We have built many regulations and many guidelines. Usually the main regulation that guides who can register .saudi; we are still restricted more so that no one can register unless he fits our criteria. And we have our requirements. We have also put a regulation that controls opening the Arabic domain name registration. And we have put regulation for submitting objection on any domain name – it's like a dispute resolution.

And we put Arabic reserved names, we built the list actually and have a procedure for that. And we published guidelines for writing Arabic domain names; we have our own rules and I will come to

that later on. And we have criteria that demonstrate reasonable relationship between the domain name and the registrant. So these are the main regulations and guidelines that we have built.

We announced our registration plan in two phases. The first phase, the sunrise phase, was started 31 of May 2010 and closed on 12 July 2010; almost one and a half months. And government entities and trademark and trade name holders can submit their registration and the name should be exactly similar to their name and their trademark. The second phase was the landrush that started in September 2010, last year and until now the registration is opened for anyone who is eligible to register a domain name.

We have done some technical enhancements. We rebuilt the registration system to support Arabic domain names and Variants. And I will speak about we have built an algorithm called Master Key Algorithm, that's all Variants in the Arabic script. And we increased of course our bandwidth because we estimated or we thought that we would have a huge landrush. However, it wasn't huge it was almost nothing.

And we installed new servers and we implemented Anycast and IPv6 to our IDN label. And we built many tools and scripts related to IDN, IDNA 2003 and 2008. When we started our registration only IDNA 2003 was implemented. IDNA 2008 still wasn't finished until last year. DNS checker and Zone editor also supports the IDN and the Zone builder under WHOIS.

We developed a registry level approach, it consists of three stages. The first stage we define the language table; the second stage we define the confusability safeguard. This is based on similarity within the language because the Arabic language has some characteristics that is unique to it.

And the third stage is the Master Key Algorithm, its script wide. The Arabic script have many characters that look alike or exactly Variant, but they are used by different languages and stage three solved this problem. The master key code will give the registrant the control to register or block a valid list of variants for any domain name.

This is the first stage – defining the language table and we used the RFC for the Arabic language that is supported by the Arab league; its' more than 20 countries around the world that use the Arabic language and there is an RFC for it, its RFC 5564. And it defines what are the acceptable code points for writing any domain name. Of course the Arabic script is a huge script so limiting the number will solve many problems.

One of them non-spacing marks, combining marks, ZWJ/ZXNJ, even digits partially – when we defined our code points we solved many problems. But still we need two extra stages. So stage number two is language confusability.

The Arabic language that we have, one letter that can be represented or written in different ways like the ʻAlif shape; there are four shapes for ʻAlif; there are two shapes for [lām]; there are

two shapes for tā and tā marbūtah. To users it's the same, but visually they are not the same so we have built this safeguard that will protect if someone registered a domain name like Akmehd without hamzah, no one else can register Akmehd with hamzah. This is the first example. The same thing for hā and 'alif, and even digit mixing. because in the Arab world we have two digits sets which is in the Unicode terms called European and Arabic index. So these two sets are used and they are known by the Arab people.

The master key algorithm I will speak about it also later on, it has a different presentation. If someone is interested there is a link at the end of this slide. It's basically, it makes the domain name safe to be used across the script. So if a user registers a domain name in a letter, this letter can be typed using Farci or Abdul or other languages. So registering a label will block the other variants automatically. And this will make the label more secure and more useable for the users also.

These are examples of the link provided, if you go to the link you will see these tools. All of these tools are published and have also a pdf describing the algorithm. We built basic rules for writing Arabic domain names. The basic rules of course Arabic labels should be at least two symbols or two characters. The Arabic label should have corresponding valid ASCII so you cannot put a string that doesn't have any representation in the ASCII label.

A-label should not start or end with a dash. And the symbol represented should match the terms and conditions specified by SaudiNIC. So we have a document that describes what are the

terms and conditions. This document lists the language table for the Arabic and several rules. These rules must be matched before registering any Arabic domain names.

The first rule is that diacritics are not allowed, which is (inaudible). So the first example you cannot register. You should remove any diacritics. The second example is no mixing between scripts, so mo mixing between Arabic and Latin. So just pure Arabic script can be used. This is the hyphen so you should have only one hyphen – and hyphen is a must for the words that if you haven't put a space between them, they will join, joining the words together will result in another meaning.

For example the last one is (inaudible) schools. If you remove this dash it will be silly orbit in the translation, so to generate a different meaning. So it's optional for the words that doesn't join together, between the words that doesn't join together, but for must for the words that if they are joined a new meaning will appear.

Digit mixing is not allowed so you can only use one set of digits, either Arabic or European. You cannot mix between them. Also, you cannot have, and this is one of the problems that even IDNA2008, the new implementation, doesn't support having numbers at the beginning of Arabic label. So if there is an Arabic domain name, numbers at the beginning is not allowed by the protocol itself. But at the end it is now allowed, or in the middle. So you can see of all the examples here the acceptable example is the last one. And if you just register it with one set you can also enable the second set of numbers.

Okay, this is the safeguard that we talked about. If you register a label the variants will be, you can enable it just no one else can. Just only the registrant for the main string can use it. So this is the concept actually for the master key. So if you have the registered label you have the master key for it. No variants can be registered for other except for the initial registrant who registered the master key. We have variants within the script. So if someone registered a domain name using eh Arabic keyboard and he wants users across the globe or from Pakistan to reach the domain name, he can use the older language and older clef, so he can enable that variant across the script. And this is guaranteed by the master key algorithm.

For awareness and support we have published videos on YouTube and guidelines and there are many pages, many frequently asked questions, we have done training, we have done presentations just to help users know about Arabic domain names, how to hose them. And this was one of the obstacles actually I faced, and I will speak about it at the end. We have a blog also for the Arabic domain names just to share experience and to help user to know what is the Arabic domain name.

What is next for .saudi? We have one main issue actually for that, which is IDN TLD Variants. Let's assume that someone opened the newspaper or go to the street and he sees a sign saying visit Mecca (inaudible) in Arabic. If this user in Saudi Arabia he will use Arabic keyboard. If this user is in Iran he will use the Persian

keyboard. If this user is in Pakistan he will use Urdu – all of them within the Arabic script.

So the same label, if he tried to type it using his keyboard, there are different representation for some of the letters. And you can see here the first string in red, Mecca in Arabic keyboard and (inaudible) here, and Mecca here on the Urdu keyboard – the strings are the same but the key and the code points are different. So now any user in Pakistan, if he tried to type this domain, he will get DNS error. Why? Because that key he had typed in his keyboard has different code points. And to solve this issue actually we need to have a TLD Variants.

In our registry we already have the master key algorithm that will help us so that Mecca will work, but for (inaudible) it will not work. And this is one of the problems that we are facing so actually we need, hopefully ICANN has announced IDN Variant Working Groups and we are participating in the Arabic working group.

IDN support, of course one of the problems that we have faced that only web is working, email is not working for IDN. So, only websites are working but when it comes to emails, still, the standards haven't finalized yet. Even search engines have problems when you try to, for example, "site;" in Google just to make sure you are searching within the site, it's worked with Google, with Bing it doesn't work.

Even the new applications have lots of issues and whenever they announced the new version we directly go and test it to make sure that it is displaying the labels correctly and we already participating in the IDN application that's headed by VeriSign. We attended two meetings. This meeting we couldn't make it because it was difficult for us or they announced the meeting late and we haven't done the preparation to attend it.

So the most problems is the right to left, this is one of the uniques for the Arabic scripts. So right to left in some applications is not displayed correctly and sometimes they don't convert, they display the A-label. Still some of the applications don't display the U-label they display the A-label.

Lessons learned – TLD variants is a must have in order for people to reach the domain name from different places around the world without any problems. And this should be transparent to the user. We have faced that hosting companies, both local and global, does not support new IDN TLDs. Especially we have done lots of communication regarding this issue, even with our cloud commuting vendors. They are a new technology, they should adopt even the IDN, the IDN is new. They haven't adopted it yet and they said they need time till we adopt it.

We need to coordinate it with our application provider to enable it. And this was because our registration systems require the domain name to be hosted before you can register it. And this was one of the main obstacles for our clients. They said they go to the hosting company and they said no the TLD is not recognized. And we

asked them to find another hosting company and they have did that actually.

One of the things also, hosting vendors like (inaudible) usually add ASCII. So www in ASCII then dot Arabic domain names and this is so (inaudible), so hosting providers of even application providers, they need to do some work there and we are communicating and coordinating with them. Also we are testing any application to make sure that the right left problem, there is no problem with the right to left display. And marketing is a must, we haven't done any marketing and that's why our number is very low. And also there is another reason that we are having a very restricted policy. Hopefully when we open the registry-registrar the numbers will go up.

That's all and thank you. For more information I have posted our links for Arabic and for ASCII websites. Thank you very much.

Eberhard Lisse: Thank you very much for one of the more interesting presentations I've seen here in quite a while. Because I had no clue how complicated this actually is, but then if you do something Fast Track of course you don't weight the implications. And I mean this idea with the space in between, that is something that I would have never, never thought of. And it probably only came out in practice, or did you envision this right from the beginning? Did you anticipate it from the beginning or did you only figure it out during implementation?

Raed Al-Fayez: Largely we have done lots of research and studies, but we have participated in the Arabic domain name pilot project, so we were opening registrations as a test bed for about three years and the numbers were low also. So we had a sense. But we expect when it comes we expected that the application will be ready because we have communicated with them more than two or three years ago regarding issues in the display from right to left.

But still the IDN solution is in the application. It's not something in the DNS itself. And this will put burden on application vendors to make sure that they build the solution in the correct way. Actually they built it correctly for left to right, but right to left is not the case. Sometimes there are problems.

Eberhard Lisse: How many, my last question, how many names have you at the moment registered in Arabic script?

Raed Al-Fayez: We have around 2000 domain names and 300 variants. But for our registry we have a total of 24,000 so our number even for the ASCII is not high. And the main reason, of course I have shared with you the lessons we have learned and the obstacles, and I just want to let you know that our regulation is restrictive mode, no one can just come and register.

The names should be exactly your official name or we have a document that gives you what names that you can register. Official name of application or it needs to have something related to the registrant. It's not something that you can go and register whatever names you want. And by the way it's free of charge, but we require the domain name to be hosted on DNS servers, on at least two DNS servers.

Eberhard Lisse: My advice is to at least charge them $1 or whatever so that you send them an invoice once a year. Then you will know whether they're still alive. Because in China they had this issue, they made it free and then they had 50 million or something registrations and 49 million or so had to be cancelled and they had to hire 2000 people to sort of do this manual. We have 2500 names in but we learned it early on.

Even if you charge them $1 then they will get an invoice and they'll say wait a minute do we really need this and they will pay the dollar if they're alive and if they don't pay, you remove the domain and you get rid of the dead wood. Any questions? Oh yes, I forgot we must always include the remote audience.

Julie Hedlund: Thank you. We have a question from Martin and he thanks you very much for the presentation and he wants to know what the equivalent would be to www and how the requirement for applications code is showed up in open source code.

Raed Al-Fayez: What was the second question again?

Julie Hedlund: I'm not sure I understand this correctly but it says literally, how this is requirement for application code showed open source code?

Raed Al-Fayez: I will answer the first question because we understand. For the first question what is the equivalent for www in Arabic domain names. Actually one of the recommendations for the pilot project, the pilot project started 2003 and ended 2006 or 7, and one of the recommendations was that for www there is nothing equivalent so you don't need to have www at all.

You just use (inaudible) Saudia, there is no need. Other countries use [molta] as for web; so [molta] in Arabic means site. So site is domain name. But for us, when we try to communicate or market the Arabic domain names we said there is no need for that. You don't need to have www, just put the domain name. And it's shorter actually.

Julie Hedlund: Okay, he rephrased the question so here goes. Is there any open source code that implements any of the Arabic code support today?

Raed Al-Fayez: Our master key algorithm, the registry level solution, it's already published and announced. And I believe its open source. I'm not sure. I need to get back to our colleagues. But when we built it we built it to be open source. The concept of the documents are published. For the code I am not sure. I need to get back to you regarding this issue.

Eberhard Lisse: On the question that he wanted to know what open source registry software supports IDN in Arabic script. I asked Andre, Fred apparently has never been tested on that yet. I'm sure it will support it. [.mazra] in Egypt, they use (inaudible) for their IDN. So there is one open source solution that is available but I'm quite sure that Fred, since check language is also IDN and Fred does that, I'm quite sure it can be adapted fairly easy to do it.

Roy Adams: My name is Roy Adams and I work for Nominet. First of all, Raed, brilliant presentation; really good. I think it shows complexities, again, I just want to reiterate what Eberhard said, languages can be expressed in many scripts and scripts can be expressed in many languages. And not all scripts are localized and not all languages are localized. Countries can have more languages than one. But you've added yet another dimension and that was particularly the key on the keyboards in different countries where the key is positioned in the exact same place in the exact same shape and print on it, but expressed in different

Unicode's. But there is a question as well – the question is have you thought about provisioning similar strings by using things like d name or c name on the service side?

Raed Al-Fayez: For d name, it doesn't work. All of us know the problem with d name and you don't inherit data codes…But for what we are doing now is somebody registers a name and he asks for the variant, we give the variant to him. So if somebody registers a domain name we have a key function that will generate a master key and all variants will be blocked from other registrants. He will be the only one who can enable one of the variants.

And we he asks for the variant we double check that the requesting variant is exactly representing his name; it doesn't represent another one. And all the cases that we have received all our legitimate and they represent the same thing. So that key will be added in the registry database to protect. So WHOIS lookup, the search, you will search for the string and for the key. So this is the new dimension that we have added to our registry. I hope I have answered your question.

Eberhard Lisse: Okay one more question because we're running a little bit late. But I find it so interesting that I'm not going to cut this off.

Stephen Deerhake: Stephen Deerhake, AS Domain Registry. Excellent presentation. Can you give a quick rundown on the policy issues surrounding the TLD Variants, it seems to me that ought to be pretty open and shut, but it looks like it's been going on for a while.

Raed Al-Fayez: Which policy?

Stephen Deerhake: The TLD Variant.

Raed Al-Fayez: The TLD…

Stephen Deerhake: You're looking for, besides your one Arabic script to the right of the dot you're looking for character set variants.

Raed Al-Fayez: Yes. The slide, are you asking me to show the slide or what exactly?

Stephen Deerhake: No. I was curious as to what's the holdup in getting these variants in the root?

Raed Al-Fayez: Okay. When we submit our application to ICANN we put our main string and we put variants. ICANN gave us the delegation only for the main string, but the variants still I don't know what's ICANN thinking, but I believe they want to know is there a way that they can make sure that variants are variants and not something else.

I don't know why, but they asked an expert in the Arab language at least who said yes this is a variant. Even some of the Arab countries register their name with a variant, not the base string, so they switch it. Because many users write the name of the country and that way it's not the official way. And one of the things that 'alif in Arabic language, usually you can write it with hamzah above or below, or sometimes you put the 'alif and everyone will know that this is 'alif. So this is a common use of that letter.

But still I think ICANN is still trying to find out a solution how to define variant, how to adopt variant, then what are the technical implementation. Are they going to delegate the variants or have d name or something? I'm not sure about what ICANN will do but we are waiting for them actually. We are waiting for them. Our registry has the solution for the variants.

We cannot implement it because our name Saudia and (inaudible) can be written in three different ways depending on the keyboard. So this is a problem. One way now it's working so any user has Arabic keyboard around the world can reach Arabic domain names, other users will use Persian or Urdu or other keyboards, the domain name, our label will not appear to them.

Eberhard Lisse: The usual thing is back to the drawing board. Thank you very much for this fascinating presentation. Thank you so much. The next presentation is Janna Lam from IP Mirror, will talk about IDN from registrar perspective. IP Mirror is a Singaporean registrar. She will tell us a little bit more about it. I know her because she registers domain names with us using EPP so when I first saw Singapore I said she's the one that I have to ask to give us a little presentation.

Janna Lam: Good morning. Thank you for the chance to do this presentation. And I would also like to just as our SA registry has just finished his presentation, it's a continuity to share my experience as a registrar to actually share this IDN experience we had. Just a brief introduction about IP Mirror; we are actually a registrar in many ccTLDs in the world. We have over 100 accreditations around the world. And of course, having so much accreditation, when there is an IDN launch we have to do it and I want to share with you how we process it and what are the challenges we face.

So for a brief introduction, probably just take ourselves back through the history of what IDN Fast Track process. It was first approved by ICANN on the 22nd of April and I think just a few countries like Saudi Arabia, United Arabic Emirates, the Russian Federation, and Egypt were the first. And of course they are actually the first to be actually entered into the root servers on the

5th of May. Just to give you a brief introduction of how many IDNs we actually have participated. In Asia-Pacific itself we have done Taiwan, of course it's .taiwan in Chinese; .hongkong; Korea, which is (inaudible); Thai, which I do not know how to read these characters; Singapore and of course in both Chinese and Tamil; Sri Lanka, New Zealand, Malaysia and Vietnam. A few of these is actually now in the process, like Malaysia, Jawi; Singapore is especially launching soon; and a couple of these have been handled in the past few months.

So, in the Middle East these are the few extensions that we have handled. So for Israel itself in Hebrew, it's actually not .israel, but it's the IDN.co.il. So the IDN is part of the domain name but not the extension. And of course in the gTLD level, we did .org in the (inaudible) Asia, Chinese, Japanese and Korean. I think .asia actually launched only these three languages for the time being, and others like Estonia, Slovenia and Ukraine.

Just to give you a brief outline of what kind of implementation is all this, just to share with you a combination not in specifics TLD. Usually the ccTLDs were launched at a government phase for sunrise; that means priority is given to the government. And usually as registrars we do not get ourselves involved in this phase. Sunrise for trademark, priorities is always given for trademark holders, just like Saudi Arabia as we have mentioned. And the trademark has to be in Arabic script

Sunrise priority given to the existing IDN. For example, IDN.sg. If you have actually registered an IDN in .sg you'll probably be

able to get a priority for the same IDN but .singapore. That is the extension that is going to be launched. And of course there is the grandfather phase, which Taiwan and Hong Kong actually use this as part of the priority. The correspondent IDN domain and the .tw and .hk automatically gets reserved. We call this bundling and the other process would be for korea.(inaudible) and IDN in Estonia.

Now the complication in these two countries is that they require us to submit documents via EPP, so and that is another area I think through this implementation, just as complicated as what our SA Registry has mentioned. In Asia we have Chinese characters which have traditional and simplified versions. So the traditional and simplified versions is causing that variant complication as well.

And we faced this problem when we were launching .asia IDN. So these are the different phases that they actually have implemented; registered trademark in any country; registered entity – either the exact match of the domain name, of the entity name located in the CD country. The CD country means countries that fall into the community of Asia under the ICANN definition. Registered ASCII .asia, which is if you have for example, if I can translate IP Mirror into Chinese I will be able to get the Chinese characters of IP Mirror.aso in that sense, in Asia.

Registered IDN in any ccTLD located in the CD countries. So of course because they are only launching Chinese, Japanese, and Korean so if you have an existing Korean or Chinese domain names already in these countries you actually can claim priorities

to it. Of course another one that's more complicated is the extended protection. So what .asia has actually tried to launch is to combine trademarks with anyone. Take IP Mirror's domain as an example. We have a trademark for IP Mirror and because we are in the domain name business I can combine IP Mirror's trademark and say IPMirror.domain and file for this application. This is what we call the extended protection.

Where is the complication? It's because they are launching all the phases at one goal, so they are not saying "Okay, from January to March is phase 1; from April to June is phase 2." They are launching these five phases at one goal, so this is the complication.

Now to share a little bit of experience with you about the problems in the implementation, they all have different rules and regulations in all the different phases as I explained earlier. So the collection of different types of documents and information – during the trademark phase we have to collect trademark information, and I mentioned before countries like Estonia and Korea actually want us to submit documents. So collection of documents from the clients to be submitted to the registries is necessary.

Control timing in between the phases – so they all have different launch times, and there's cutoff times to certain periods. So this is one of them. So how do we handle all these implementations? We have a system called the ccTLD Box, which we register all the domain names in the world. So this is the Box.

Now just to give you a highlight of how we handle the IDNs during this phase, and just to highlight some of the features of the ccTLD Box that we have: in our system we are able to configure the different phase in line with the policy, like trademark phase, the locality restriction. Like in Saudi Arabia, the only organization in Saudi Arabia is allowed to register; so it is for Korea, and also certain TLDs in Hong Kong; and of course the unique requirements such as document submission.

So we are able to configure our system to accept all these different rules and regulations. The ability to set the timing for each phase. Imagine if the phase is actually open at about 2:00 pm Singapore time on the other side of the world; we will be able to set the timing so that we still can have a good night's sleep. And of course the ability to configure the different language types, just like .asia. They are not launching all the languages but only the three major ones in Asia – Chinese, Korean and Japanese. So we are able to set the system to only accept these three language types at the moment but in the future, if .asia is going to launch more language types we can add on at a different phase.

Ability to set minimum and maximum characters – it's very different from ASCII because I think IDNs do set a limit to certain characters. So we are able to actually set the limit, the so-called limitations of the number of characters you can set in the system. And there are actually more features we have which it will be never-ending. This session head is going to actually chase me out

if I go on too long but basically we make our system in such a way whereby it is going to have ease of implementations for IDNs.

We recognize that being a ccTLD registrar, there is a lot of different rules and regulations that we have to cater for, and that is why our system is robust in the sense to be able to handle all these different rules and regulations. So just a quick overview of the system: we are IDN ready. We are able to support any languages, so the cost of ownership is actually very low because it's on a hosted platform. We do provide this system to our partners who work with us, and if they were to use this box to register IDNs or any other domain names the cost is very low; reason being because we have built the system in such a way whereby the business rules can be easily configurable.

The maintenance is low and we do updates once a week, so we get four times a month of free updates. Scalability – because of the configurable business rules we have we will be able to actually scale the system to any level. High availability – of course, based on the low (inaudible) and high availability, the backend system that we have. Payment gateway – the system is able to actually accept PayPal, WorldPay, AsiaPay, AnyPay. AnyPay incidentally is the payment gateway in Korea which you have to use only in that country. And of course AliPay is for China and for many more.

Configurable email templates – you can add your own. Multicurrency, which of course we support any currency in the world. So when there is a promotion campaign, the system is also

configurable to that extent. As many of the industry players here probably will understand we mentioned about variants, we mentioned about bundling. So if a registry decides that they want to offer a variant with a fee or offer a variant for free, the system is actually able to cater for that. So and of course this is one system for registries and registrars.

So my conclusion is registries show a lot of creativity. Through all of these launches we realize that the registries really come out with many, many different types of launch patterns, and to be honest I think .asia has given us the biggest headache with so many launches at one go, but this shows how every country is very different. So to be able to cater for that is very different.

Language barrier: of course coming from Asia we are comfortable with the Asian perspective, but after hearing our speaker from Saudi Arabia we realized… I mean I do not recognize even a single character up there, but isn't it interesting to know how it works? And to us it's a very challenging process. Ever since the IDN Fast Track has gone through, the introduction of IDN has been really busy for us. We thought we could have a rest a little bit but it doesn't seem so, and it has been very challenging. And seeing the trends of how IDN has gained popularity is kind of encouraging.

Of course usability: we did mention about it. Hosting companies are posing a problem to hosting IDNs. I think it's the same everywhere. IDN is really new in the industry; not many hosting companies are actually providing that service, so I believe

education is needed; and of course to do this, to make IDN a successful launch. Also localization, because many of these IDNs are meant for local markets. It's just like the Chinese character is for the Chinese community. If you look at Asia, the Japanese character is for the Japanese, or Korean, the Hangul for the Korean market. So it's very localized and I believe this will actually help people who have a language problem in those local markets.

So why IDN? In our experience when we are doing all these launches, we believe that more education is needed to educate on the use of IDN. As I mentioned before it's not just to the end user but also service providers. That's the end of my presentation, thank you.

Actually, before I end this IP Mirror has a booth and we printed a small little booklet in English. As you know probably, you might have heard how Singapore speaks a varied mixture of English. In this book we have IDN Introduction in Chinese, in Korean and in Japanese, and also some local food introduction. So I encourage everybody to visit our booth, get a copy of this and understand the localized language in Singapore.

Eberhard Lisse: Thank you very much.

[Applause]

Eberhard Lisse: Coming from a place where we speak [Nomlish] I appreciate the English to some extent. Any questions?

No questions? Okay, then you are released and you may step down as they say. Thank you very much.

The next speaker who is just setting up his line – Tan – he's from SGNIC, and as it has become sort of a tradition we always have the host give us a host presentation. He will speak about DNS Resilience, which I think is also a good idea. I've also asked him to say a few words about how they run their registry so that we just get a little bit of insight there. There you go.

Ryan Tan: Good morning everybody, my friends. On behalf of the Infocomm Development Authority of Singapore we welcome you again to the (inaudible) Island Country. My name is Ryan and this morning I hope to share some experiences that we have done to improve DNS resiliency in the .sg domain space. But before that, just a background of SGNIC.

SGNIC is the internationalized domain name registry for Singapore. As you can see on the map we are a very, very tiny country on the tip of Malaysia. We have a population of about 5 million. I'll give you an idea of where our Infocomm status is at. We have our 8 terabytes of summary cable capacity. Broadband penetration, this is the household broadband penetration rate is about 81%. The mobile penetration is 145% and I was told that about two-thirds are using smart phones.

We have quite affordable 3G plans, usually bundled with the iPhone, and we have about 7,500 hot spots around the island that you can go on wirelessly. So you can see all the cities are actually quite online. That gives us a lot of pressure as the .sg registry to keep the domain names alive. In terms of numbers we are quite small actually, 329,000 domain names. Most of it goes in the .com.sg space, and the second most popular category is the .sg space.

We are actually, legally we are set up as a Private Limited Entity. We have a board of directors and a management team, and although we are Private Limited we are 100% owned by the Infocomm Development Authority. But in our day-to-day workings, they only interact with us if we are deciding policies of national impact. Otherwise they leave it to us to decide the other policies.

We run a shared registry/registrar system, and we also interact frequently with external organizations like ICANN, IANA, AP TLD, ccTLDs, AP NIC and the local CERT, (inaudible) CERT.

Well because we are sort of linked to the government, so we try our best to adhere to the very strict e-Government Security procedures and controls. And we are able to leverage on the government's secure infrastructure, and (inaudible) for example, we're able to host our systems in the government data center and we're able to subscribe to a cyber-watch center service which is essentially an IDS service.

And this is quite interesting. The government has a directive to direct the IPS to flush DNS in case of emergencies, so we have a special link to the ISPs to say "Hey, if there's any problem, we want you to flush your DNS. You have to do it within…" I think it's about four hours. And as we wish we frequently share and collaborate with various stakeholders and on (inaudible) processors and security with a lot of findings, like what we are doing here for instance, at this forum.

For our DNS infrastructure we basically have four things that we look at. Number one is that the hardware and software systems, we always try to design them in such a way that there's no single points of failure. For the network, we make sure that the network is multi-home and that it's protected by the IDS and the IPS. And all these systems are placed in a physical location in a highly secure data center which as I mentioned earlier is the government data center. They are so secure that not everybody can go into the center without being checked by the department known as the Internal Security Department, much like a "secret police." So they do a background check on who you are before they let you into the place; even the person who delivers the [hot dogs] to the center is being checked

And of course we need to have separate [DL] sites. All these are being controlled by the IT control procedures that we have and monitored continuously on a 24-hour basis on a recently set-up entity known as the Business Intelligence and Response Center.

For our name server constellation, this is pretty much quite standard across many ccTLDs. We have a database, a stealth DNS – takes the data across to its own generation. It goes into a primary name servers; in our case it's all (inaudible). It's not any public traffic. And then it gets propagated through [C-sig] protected by firewall to our secondary DNS servers which are anycasted. So anycast is our solution for the DDOS resiliency approach.

So at this point I wanted to ask you a question. As a ccTLD manager what is your worst nightmare? I don't know but if you ask my worst nightmare it's to wake up one day and somebody tells me that your entire nation's domain name is taken off the internet because of a problem with your zone file. And it gets worse when people discover that… If people don't know that you have a problem and you fix it quickly, that's not a big issue. The biggest issue is your Prime Minister, your Minister, even the grandmother next door knows that you are responsible for keeping the entire nation off the internet.

So we're very careful with this particular portion. The next slide will show you how we are dealing with our zone generation process. So there are multiple safeguards. First of all, before we reenter we do a lot of careful reading and writing error protections for data from the database. For example, we will not read the database in the one chunk, one big large order with entries all together. We read in chunks of say 5000 records and then move on to the next 5000. This will ensure that we have a high success rate

in reading the data; then for us in writing we make sure the data is being returned, that there's no disc errors and stuff like that.

And then if you look at the pink-colored section we have a "white list" protection. What is this? There are certain VIP domain names that can never be off the internet even though they forget to renew. So this is where we put a list, a white list to make sure that we protect some of these very, very important domain names that can be never off the internet for a lot of reasons. Of course SGNIC.sg is one of these white list protections.

After that we do a lot of syntax checks, and then we do a zone size comparison check. Now, this zone size comparison check is very important. Just as you can imagine your domain name, you have 1000 domain names and you do zone generations every four hours. So between each zone generation you will expect that the zone difference cannot be more than say 1%. So as a result of your zone generation process, if your new zone file is 1% more or less – then your last generation, something might be wrong. So you want to be a bit more careful; you pause the process and we do a manual intervention to make sure that the records are still correct before we allow the zone gen to proceed.

And then we load that zone file into the BIND to make sure it loads and after it loads we are still quite worried. We are thinking that maybe the name servers that we grabbed from the database is wrong, so we have an external system to compare our database records versus what is published via the BIND.

So as you can see we have quite a complicated zone generation process. Thus far he has worked quite well for us, but the thing is we are still worried. So anytime for example if you have some code changes or zone generation scripts something may happen – maybe there are some problems with the machine that's performing the zone gen. So as a result the zone is still taken off the internet, so what do you do?

So we're thinking if that process does happen to us we do not have a lot of time to go in and figure out what's wrong with the zone generation scripts, what's wrong with the server and stuff like that because the entire nation is putting a lot of pressure on you to get the problem fixed. So what we have is we have emergency link server. On a daily basis we will grab a copy of a good zone file that we have, we know is working and chuck it into the emergency link server. So when we have unforeseen failures in the zone generation logging we will do what is called a rapid IP swing to the secondary servers.

Now this process has some problems. Because it's an old copy of the zone file you will find that maybe 1% of the changes made of the new domain names being added is not inside, but we're thinking this is an okay compromise to make 99% of people happy and then deal with the rest of the 1%.

So as a conclusion I just want to show you three points that we learned so far, that I think we think that collaboration among the governments, ccTLDs and especially the ISPs is very important. We have some processes in place that I mentioned earlier to get

them to flush the DNS in case of emergency. We also think that the collaboration and information sharing among the industry stakeholders on best practices and security alert findings, such as Conficker worm mitigation is in the way we went through two years ago.

The last point I want to make is of course equally important, that we must equip our DNS infrastructure with robust configuration I think of all the possible problems that could happen. So with that I thank you. Is there any questions for me?

Eberhard Lisse: Thank you very much. I have one question from the Chair as usual. Have you ever been phoned at home at night by the Prime Minister – "Wait a minute, what's happening?" Have you ever had a catastrophic failure that you're so worried about?

Ryan Tan: Officially no. That's a good answer, right?

[Applause]

Ryan Tan: Between the guys in the room we've had partial failures. It's not the entire zone; a certain portion of the name server is gone. We are lucky. I believe most countries have such problems. I think Sweden was the only one who was daring enough to admit about

an incident and publish a report on it. Some of us as I know for instance have signed [CDC] and then…

Eberhard Lisse: The Germans also admit it.

Ryan Tan: Who?

Eberhard Lisse: The Germans also admitted it.

Ryan Tan: Did they? Good for them.

Eberhard Lisse: I mean for me, running a very small zone, I upload it hourly and I keep a timestamp in the filename and at an external server. And should we lose communication the operator, [IOC] operator, they will figure it out – Peter is sitting there – and then they can manually load the last one that works until we straighten out the casualties. I think it is a good approach to have a backup in place and it's better to worry about 1% than about 99%. Anyway, any questions from the floor? No questions, thank you very much.

Ryan Tan: Thank you for the time.

[Applause]

Eberhard Lisse: So João is going to be giving us a webinar offline. ISC has been doing a webinar for DNSSEC and when I saw this about, what was it, three weeks, four weeks ago travelling I thought "Wait a minute – that's a good thing that we can repeat," not necessarily to preempt Simon's DNSSEC for Everybody but I missed the presentation or the webinar because I was busy with delivering somebody. So now I thought we'll have a look at it here, and that small talk is getting you ready. You're online.

João Damas: Hi, so I'm João Damas. I work for ISC. You cannot hear me? Let's see, is this better? As I mentioned I work for ISC as Eberhard just said, and indeed we are having these webinars where we go and try to explain to anyone who wants to listen different aspects of the DNS and how things operate. And obviously since we are ISC we do tend to have a focus on BIND and how to use BIND to make things work for you.

What I'm going to do here is not actually one of the webinars. The webinars are available online – all the past ones are archived. You can download the slides and actually even the recording from the ISC website, www.isc.org/webinars. There's also there a list of the upcoming ones, and in trying to cope with the fact that the Earth is

still a round planet and we all have different time zones we usually, when we have a webinar we have it at two times during the same day to try to make it available to people in different time zones. So that's just a note.

What I'm going to give here is not a lot of technical detail but more a little bit of an overview of the ideas of how this whole thing with DNSSEC works, and in particular how you have to think about the keys. So what you can see there is just an outline of the things I intend to cover.

Starting with DNS and keys. Nowadays we tend to think about keys only involving DNSSEC but that's not quite true. There have been keys in the DNS for a long time, even if you don't think about them a keys, used in things like TSIG. Now, TSIG uses a type of cryptography that's called shared secret cryptography and it's good, it's lightweight, it's fast to process. It has one problem that has to do with the word "shared" there. It is the type of key where everyone has access to what would be called the private key in the normal cryptography, so everyone who wants to use that has access to the secret.

And the problem with secrets that are shared is if they are shared too much they stop being secrets, so these types of keys are useful for securing communication between servers which are a small set of involved parties, but do not provide good scaling properties; does not have good scaling properties to be used on the internet at large as a means of securing the DNS. And that's why this additional thing that's called DNSSEC was created, and it uses

public key cryptography which has all the properties that you are familiar with and are used in things like securing the web or PGP or so many other applications that we use every day.

So DNSSEC keys, which are one of the fundamental elements if you are talking about DNSSEC key management which is supposedly the objective of this talk, are basically normal public cryptography keys. They can be using any number of well-known algorithms out there. Typically the most common one is ISA which everyone is familiar with. It's the one that for instance is used in DLS for our SSL for securing web communications, but there are other options. There are options like DSA, there are options that have a more original focus than just the ghost protocols which were defined by the Russian cryptographers and are mostly used in that region of course.

A typical DNS key as representing DNS has data about the key obviously, and it has a few other things attached to it. One of them, because this is DNS after all is the name – the name of a domain name where this key is meant to be used. And then there are some little bytes here and there that identify protocols and so on. I'm not going to spend much time looking at it; there's one example of a key later on we can point things out.

When you come across DNSSEC keys you very quickly come across two different concepts, which are the key signing key and the zone signing key. These are actually the same. The distinction one chooses to make in this context is mostly something you do for operational convenience. It doesn't say anything about the

property of the keys themselves – they are exactly the same. There is one little thing that changes that is mostly a signal to the consumer of the information; it doesn't actually change what the (inaudible) means.

Because DNS is a protocol that has a limited amount of signaling bits going back and forth we also chose to overload some fields in the key to do signaling, and so for instance there are different ways of signing a zone to provide authenticity to non-existent responses – what's usually called NSEC versus NSEC3. I'm not going to describe those here because that's nothing to do with how you manage the keys, but those actually, whether a zone uses one or the other is something that gets signaled through the type of algorithm in the key.

So key signing keys, zone signing keys. When you have – and I put there the two keys that I currently use at isc.org the domain – the only distinction as I was saying apart from the data, the keys themselves are different even if the algorithms being used are the same. The distinction between zone signing keys and key signing keys is defined by the number that's circled. A key signing key has one extra bit set to 1 that changes the value that normally is 256 to 257. This extra bit is called a secure entry point big – SEP – which is kind of a somewhat strange name to have now that DNSSEC is building a chain of trust, but it is meant to be the key that you use to verify everything else that's in the zone, including other keys.

And the zone signing key is the other one – the one that doesn't have that bit set. If you have used these in software in your DNS zones you'll soon figure out that the distinction is mostly for humans, that the software itself doesn't actually care whether the bit is set or not. It uses it as a hint to facilitate search and to try some keys before others but it doesn't actually discard any keys just because the bit is set or is not set. I will explain later a little bit of why this distinction is made and what is the operational advantage of having that in your DNS.

The next element of DNSSEC is what you do with the keys, and the keys of course are used to sign things. How you use the keys in this context of producing signatures is basically as is outlined there. What most people do out there is that they use the key signing key – the KSK – to sign only the set of keys in the zone, and they use the zone signing key to sign all the other contents of the zone, the zone itself. And hence that's where they derived their names from, right, because the key signing key is only used to sign keys and the zone signing key is the one used to sign the whole zone. And of course signatures are what actually provide the means to verify that the data has not been altered since it left the authority servers and it reached the clients.

There is this third element which is necessary to put everything together in working order, and this is an invention that wasn't around when DNSSEC was originally defined more than ten years ago, but that need to be added so that things made sense to people who had to operate it, and this is the delegation signer record.

What this record does basically is provide a link between zones in a secure manner without imposing too much of a burden on the operators of both the child and the parent zone in keeping this information up to date. What these delegation signer records are made of are basically a hash of the name of the domain that's being secured and the key data.

These are put together and use one of the traditional hashing algorithms out there – SHA1, SHA2, whatever your choice is; nowadays probably SHA2 since changing algorithms is a pain so why not choose a good one to begin with, for instance like the root has done – and the few additional bits to identify what protocols are used, which are basically hints to humans and software. One big distinction about, one big point about these records which look a lot like NS records in delegations, somehow they are not delegating the DNS information but they are delegating the security information, so they have somewhat parallel uses.

But they have one big difference – when you have a delegation DNS you release the name servers both at the parent zone and at the child zone, but it is clear from the DNS standards from the very beginning that the records at the parent are mainly there as a hint so that you can continue your search. But the authoritative ones, the real ones that you should believe are the name servers that are present on the child zone because that's the zone whose services is providing the ID service. So the child is the one that knows the name servers for sure.

The child better make sure that the parent knows at least some of the NS records properly otherwise it's going to be very hard to find that zone, but when it comes to delegation signer records, the DS records, this is completely different. The parent is actually the one that has the authoritative information because the parent is the one that's going to sign the record and provide the security link downwards.

So what the DS record does is enable us to provide a chain of trust where hopefully starting at the root you can go down the tree of the DNS and verify the links of security between each nodule until you reach the one that you want. However, the DS record is as I said before produced from the domain name, which everyone knows, and the key data, which again should be down to the child zone, the one that's been using those keys to sign data to provide.

There is still I think at this time a lot of discussion about who should be generating the DS records – should it be the child zone administrator that produces these records and then sends them up to the parent zone for inclusion, usually a registry of some sort? Or should the child communicate the data for the key and should then the parent be the one generating the DS record and including it? This mainly has to do with which choices of algorithms are used by each of the two zones, and I think this discussion is going to be around with us for quite some years, and different registries are opting for different solutions to this mechanism.

The key that you use to generate the DS record is usually the key signing key and not the zone signing key if you are using such a

distinction. We'll see why in a minute. When you put all these elements together what you get is something that looks like the DNS tree but is actually somewhat of a parallel tree because it's not talking about delegations and how the DNS links itself together – it's talking about how the security of the DNS links itself together.

So today we have the root zone which is signed and hopefully you are all aware of that and have a copy of the root key in your possession, verified hopefully; and as you are beginning to see more and more of the TLDs, the next level down that have been signed are including their DNS records in the root zone. Hopefully more and more will add it – I think we are up to 50 or 60 now so there is still some way to go, but it is progressing at a very decent pace.

Each TLD will have or will eventually enable signed zones underneath it and will use its own set of DS records to link down from the TLD to the SLDs and so forth until it reaches the end domain that you are looking for. Some TLDs are not signed or don't have the S records in the root zone, and these create what's called islands of security.

If you want to reach those islands of security in your resolvers you'll need to have access to the keys that are used by those domains and in some other parallel way that's not just progressing from the root downwards, because there are points where this link will be missing. And it is what are usually called [thrust anchors] or security points, hence the name of the bit in the KSK that we

mentioned before. All the resolvers that I'm aware of that support DNSSEC capabilities are able to configure more than one thrust anchor and so this is just common operational procedure.

When you have keys in a zone or when you manage keys there's always the question of how long should you keep using any given key. In the DNS this is particularly so because keys have no expiration dates; keys by themselves will remain valid for eternity. This is different than signatures. Signatures that are used in the DNS do have expiration dates, and you have to decide what those dates will be to begin with when you sign the zone using your keys.

And in fact you have to be careful about those dates because it's been one of the most common mishaps in secure zones that people forget that these signatures expire, and they have to be renewed; and if they aren't renewed in time before expiration then you come across problems where suddenly your name, your domain names go insecure and that's not fun. Still, as I said, it has been one of the most common issues because the software as it was originally published when we had no real operational experience in the field didn't help you a lot in keeping things running. It did the bare basics.

Another thing that can happen is that algorithms that we use today that are believed to be strong may not be strong forever. I mean computers keep increasing in their power so brute force attacks become more plausible as time goes by. So that is usually solved by increasing the size of the key, but sometimes there are

fundamental things that are discovered by algorithms that allow you to factor the keys and then you have to really look at changing algorithms and therefore changing the keys.

It's possible; it's not probable. The third one, the third item down – unauthorized access to the private key – is probably one of the most probably scenarios of why you would have to change the keys, because in the end you can have all the automation that you want but there are still people involved in this and accidents can happen, do happen. When you have to change one of the keys is when you find out how convenient you might be to have this separation of concepts between the zone signing key and the key signing key; because if you do have these splitting personalities, even if the keys are all the same to the computers, when you have this split changing a zone signing key doesn't imply any change outside of your zone, outside of your domain of control; whereas changing a key signing key implies that you have to generate an (inaudible) record and you have to communicate with your parent zone.

It no longer involves only you and your operation; it starts involving someone else's operation as well. You have to communicate this change upwards to your parent zone. So having the split, while it provides, it enables you to have a very big – bigger than normally used – key that is well protected, hidden away somewhere with very limited access that you only touch when you need to sign the rest of the keys when you change the key set – and because that's something that only the people who

are operating the domain itself need access to it's easier to control. And you can then use more lightweight keys and have less protection, hopefully that doesn't mean no protection, for the other keys – the keys that are more disposable and used for everyday operations.

What I mean, well I'll talk about that later. When it comes to rollover if you really have to do it, and there are many questions about when you do have to do a rollover or you just need to be prepared to do a rollover, because keys to my knowledge, the keys that are used in the DNS have actually never been broken by cryptanalysis or brute force attacks so far. But still there remains the issue of how much trust you put in a key that has been used for prolonged periods of time.

So even if you don't think that a rollover might be necessary ever it might be a good idea to have some way, some documented way of actually running the keys in case an accident happens. And the problem you are all familiar with with operational procedures that are not used is that the day that you use them is the day when you find out the document is actually missing a couple of critical steps, and you really don't want to find out about that when the Prime Minister is calling you as the previous speaker was saying – that's really not the time to find out you had things missing.

And that's actually the reason why some operations do periodic, regular rollovers, just to keep in mind how the process works and make sure that everyone knows about it. It's not so much about the weaknesses in the key or attacks on the key, but more of

keeping operational experience fresh in your mind just in case you actually need it when you are into a problem.

So how do you roll over a key? And this is rolling keys that involve no change in algorithms. If you also have to change the algorithm then the steps are a little more complicated, and I'm not going to go over that because it would take too much. The details from the presentation would probably not be good enough so you would be advised to actually go fetch a proper textbook on how to do it. Also this is one of the reasons, because algorithm rollovers are harder than just simple key rollovers it's one of the reasons why you'd be recommended right now, if you are about to sign your domain, to start using something like a SHA2 algorithm rather than a SHA1 algorithm, which is older. And even though it hasn't been broken today some people say it will be broken in the next few years, so use the newer one and save yourself the hassle of having to roll it in a few years.

If you go about rolling over the keys there are two main ways of getting to do this: pre-publication of the new keys and double signing. Double signing is quite easy – you have one key, you want to introduce a new one; you publish it and you sign the zone with both. You thereby produce two sets of signatures, one with each key that are both valid at the same time. Hopefully when the older set of signatures expires you've retired the old key and keep signing with the new one.

The problem with double signing is that because you are producing two signatures and the name servers don't have a way of telling

them apart, they have to provide both signatures in every answer when DNSSEC information is requested, and so the packets, the information that's sent back to the client is double the size basically. Double the size in principle shouldn't be a problem – the DNS has many mechanisms to carry large payloads. In the practice it is a problem. It is a problem not because of the DNS but because of all of the other boxes that you find on the internet, on the networks between the DNS server and the DNS client.

Firewalls still to this date have a very limited understanding of how the DNS works and choke, stumble on big DNS packets. And some of those you don't have control over – they are simply in the path and they generate problems for everyone. So double signing is an option but beware of the problems that you may encounter. Actually I'm reading this and the beware notes on these bullets are swapped around – interesting.

Pre-publication is a different method of doing this where you introduce a new key and thereby by publishing it in the DNS you make it available to the clients, so the clients can start caching the information, but you don't use the new key to sign. What you do is you just publish it until you make sure that all the information that's current in cache is expired and once that's done you start using the new key.

It's a little bit tricky. You still have to be very careful with the TTLs and your expectations of how long things are going to be cached out there, so it puts a bigger burden on the operator of the domain to keep things working. But it doesn't duplicate, it doesn't

double the size of the response so it doesn't put any burden on the client side, on all the machines that you can come across. So it's your choice, none of them is perfect. On one of them you do more work but save your users problems; on the other side you do less work but you expose your users to a lot of potential problems. And so probably you'll have to do some research to choose one over the other.

One thing you have to keep in mind of course when talking about key management is where are keys stored, and I think one of the key questions on deciding where you're going to store the keys is that question I've put there – "For you, which is more important – the key or the zone data?" And hopefully the answer will be they are both equally important, and from that you might derive what sort of storage you need for the keys.

If your zone data is available in… By zone data I mean the database that is used to generate the zone, so your registry file, for instance, your registered database. If you have this database in a machine that's accessible from anywhere then why would the key need any more protection than that, because after all, all the key does is take the zone and sign it. So if the zone is not protected, if the zone is altered by anyone in an unauthorized way then any signature you put on it is not going to correct that problem.

All you are doing is even putting some fictitious trust in data that was bad to begin with. So this is something you have to keep up, because it's becoming fashionable to have these crypto-machines, specialized hardware and perhaps it's not necessary for everything.

There are scenarios where it's useful to have one thing but not always. And because of this, this comic strips on [XKCD] that I found very appropriate. I'll let you read it. People, particularly techies, tend to think about attacks on keys as very sophisticated things – you have to use brute force attacks by clusters of computers. And when you come down to reality things are more like what the right side says.

And this is probably one of the reasons to use HSMs, to save the poor system administrator from being hit in the head with the wrench to get the password out of him, right? So it's not so much to protect the data but to protect the people who have to handle the data.

Aside from that, well, depending on your needs you can store the key in a number of places. You can store the key in a file system. If you do so and I do that for some of my zones, for my private zones, you have to be careful of course. I mean don't do for instance what this French CA did a couple of weeks ago, where you leave your private key available on the web server and then someone makes a mistake, changes this web server configuration, makes the directories browse-able, and anyone who goes to that webpage, instead of getting the webpage gets the directory and there is the key.

If you do that of course you are dead, basically. You have to revoke the key, the key can no longer be used because everyone knows about it, everyone has its power. Cracking the passphrase that protects a private key is usually not a big deal because these

things are put in there by people, and people do have predictable ways of generating these things.

An intermediate state of key storage is to have them in the file system but have them in a machine that is offline, and this is again where the distinction between KSKs and ZSKs comes in handy. Because the KSK is used only to sign the keys you can keep it in an offline machine which you only have access to when you need to generate the new zone signing keys, which probably is not more than once a month if that. And so that machine can be kept away in a secure place. The machine itself is just using whatever the operating system provides as protections, and then what you do is you generate the keys, you sign them with the key signing key, you put them on a USB stick and take them out of this machine and go over to the machine that's actually on the network that uses the zone signing keys to sign the zone.

And because these zone signing keys can be rotated more easily than the key signing key, you have to be careful how you store them and who accesses them; but if you detect any issue with that key you can rotate it very easily because it only involves your own people.

And finally the higher standards of protection are those afforded by HSMs. Now, HSMs do have very nice features, for instance typically these machines generate the key inside them and there is no actual way of extracting the key from the machine; and if you try by opening it or doing something nasty to it it will self destruct. That itself generates a separate set of problems in that you need a

backup somehow that will enable you to regenerate that key in some other HSM, because otherwise it could be a denial of service to just walk up to the HSM and cause it to self destruct. These things self destruct on things like high temperature, high vibration; of course any attempt to open them usually makes them self destruct.

So they don't come without their own problems, but they do come with the advantage that no one's ever going to be able to access the key and copy it. For the root zone for instance those types of machines are in use because there are just too many people involved in signing the key and they really don't want any of the volunteer people that are involved to be mugged just to have access to the key.

So each of these systems has its own advantages and its own drawbacks. HSMs are not magic bullets. They provide protection against certain problems but they come with their own. Plain file system storage is easy but probably also easier to access for other people. So think about it, which is the best way for the type of zones that you are protecting.

And finally, doing this all by hand is a nightmare. When we started ISC producing BIND versions with support for DNSSEC, or at least the current version of DNSSEC it came with tools to generate keys, to sign the zone with those keys, to publish it, to do some checks but they were not integrated at all. And what we saw is that we humans tend to make mistakes and many mistakes were made.

So now the trend is that the software is moving towards more automation. I'm listing three different options that are available to anyone that will automate different degrees of the key management process and the zone signing process in an attempt to minimize the amount of errors that we have seen. And BIND versions 9.7 and after are the ones we would advise anyone that's considering DNSSEC to be using because they have features that will prevent for instance your signature from expiring. It will make sure that your signatures are renewed in time so it eliminates that problem.

It provides a mechanism for you to tell the name server when a given key should start to be used and when a given key should stop being used. It doesn't yet generate the keys themselves, you have to generate that manually, but when you do generate these keys you can provide additional information so that the server knows when it should be using a key and when it should not be using a key. So it's not complete automation; it's a lot better than what used to be. And in future versions, in particular BIND 9.9 which will come out at the end of the year, we will be also automating the generation of the new keys.

There is an Open Source toolset – ZKT – which is very lightweight, very easy to use, and if you are starting to use DNSSEC keys or if you are using them for your own zones, for instance, or even for bigger zones, enterprises and so on, it's a very useful set of tools. It automates a lot of the process, it's very simple. It doesn't do a lot of fancy things so you have less chances of errors, and it's Open Source, it's free – you can use it.

There's a third thing called OpenDNSSEC which aimed at providing all the things that BIND didn't provide when it started up and that I have mentioned – generation of keys, scheduling of keys, all things involving DNSSEC. It comes both with a key management system and something that will check that key management system – the auditor.

The only problem, at least from my point of view, this piece of software has is that it was the brainchild of a set of geeks that went wild and produced something that is really, extremely complicated. It's been simplified a little bit since the beginnings when it was basically unusable because of its complexity. I mean it had theoretically everything that you wanted to have in it but the dependencies, the consideration was such a nightmare that instead of simplifying things it would actually complicate them.

It's getting better. It's still a little bit complex; for most people I think it's a little bit too much, but here for TLDs it might just be the bill. It will again be down to you. Some people are using it in production; you are welcome to talk to them. This is software developed in conjunction by the Dutch, the Swedish and British TLDs supporting, and now NLnet Labs – all of them are here and you can approach them about the properties of this system.

Just finishing I'd like to give some thanks to Alan Clegg of ISC who was actually the author of the two webinars that Eberhard mentioned earlier and that you can download, and is an all-around good chap. And if you have any questions you can ask them now

or you can send them to that email address later on and I'll be happy to answer them.

Eberhard Lisse: Thank you very much.

[Applause]

Eberhard Lisse: I have got, with regards to figuring out whether my zone expires, Ondrej Filip runs a script that reminds me once a month, and that got to irritate me so much that we wrote a shell script that irritated me so much that I wrote a Perl script to irritate him once a month that it is not going to expire. And of course recently when I cleaned up my server I deleted the file, but fortunately I've got a backup system so I've asked the host to load the thing.

But what I'm trying to say is it's fairly the only problem that we saw was that the zone expires or this thing expires and you don't think about it, because it works, you don't need to touch it – it works on its own. But it's simple, simple Perl script which checks it and sends an email, and then you can manually redo it or if you trust your Perl script you can even automatically regenerate it. I like to do it manually.

We don't sign the zone .na with 2500 names, doesn't sign the zone with DNSSEC. Since the author is not in the room I can really talk

bad about it. I find OpenDNSSEC difficult, yeah? And when I'm connected to the mailing list I read too many emails about things needing to be fixed and bugs like this.

But it's a good concept and .sa uses it, Nominet for .uk uses it, so if you've got the-

What, they don't use it? They wrote it.

João Damas: They wrote it but they don't use it.

Eberhard Lisse: Okay, who uses it? .sa uses it, eh?

João Damas: .se uses it I believe.

[background conversation]

João Damas: Does .nl use it? No one here to answer?

Eberhard Lisse: NZ? Yes.

João Damas: Okay.

Eberhard Lisse: So some major zones are using it so it must be working, but it is not trivial to set up. Once you've got it set up and working you can forget about it I think. You don't really need to touch it unless something happens and then you have to squeal to the mailing list.

For our little thing where we haven't got a bank really interested yet, you can use the existing BIND tools, write yourself a provisioning script or edit into your provisioning script – I did that when I was sick and I was bored and had to lie in bed for six weeks.

PCH is offering a signing service. I don't recall whether you remember this?

João Damas: Yes, that was announced at the last ICANN in San Francisco. PCH and ICANN are working together to provide what's called a 'bump in the wire' kind of signer, whereby a registry would feed their servers with the unsigned data as you do now, as you publish it now; and they would take it inside the servers, publish it, sign it and then publish it out the other way, either to your servers again or a combination of their servers and your servers, or only your servers – whatever you choose.

So they would take care of all the logistics and complexity of all of the signing. It's good in the sense that you have someone to blame

if something goes wrong. You also take off a lot of load off your plate. You also give a lot of control of your zone to someone else, so it's again… There's no perfect solution – every solution has pros and cons, and it's up to you to decide which one.

But it's true this is available. I think they have a distributor site with three sites around the world: in the US West Coast, in Zurich and here in Singapore in fact is where the third site is hosted.

Eberhard Lisse: We use, I'll come to you just now. We use PCH as one of our secondaries of those so they did do anycasting for us and I find they are a reputable organization. So if you look at the level of trust, we are going to use them for the economically active second level zone, or third level.com.na. The way we are going to do it is since they are secondary for .com in any case we are going to switch them to be the primary.

So we have a hidden primary in our registry system and then we TSIG that stuff up so that there is a chain of trust from our registry into their system of sorts, and then they put it, sign it properly. And I think it's a reputable organization with sophisticated know-how and individuals, with a program backed up by ICANN – not that I like ICANN very much as you all know but as far as the level of trust goes, if a bank wants it you can show this is a well-designed system that is reputable that the banks can use.

So we are definitely going to implement this by the time of the Dakar meeting so I will report back on the findings.

João Damas: Okay. Actually I have to ask a question to Peter, who is sitting there. Is SNS going to provide this as well or not?

Peter Losher: Peter Losher, ISC. We do have plans to do bump the wire. It's being done in coordination with some of the work that we're doing for v9.9. So watch that space for the end of the year.

Kristina Nordström: Hi, I just wanted to say that anybody that goes up for a presentation, can you send me the presentation first so I can upload it into Adobe Connect? I will leave my email address up with Eberhard, that would be great so the remote participants can follow as well. Thank you.

Eberhard Lisse: It's Kristina.Nordstrom@icann.org and it's on one of the mailing lists that I sent to all participants, so all presenters should have it on their thing.

Russ Mundy: Thank you, Eberhard. This is Russ Mundy from SPARTA. I wanted to point out there's another major toolkit besides what's been mentioned today and that's DNSSEC-Tools and it's available at DNSSEC-Tools.org. It's all Open Source Berkley license and it was created initially to help out the earlier versions of BIND that

were really hard to do DNSSEC with. All of that capability's still there. There's a batch of applications that we've also worked on but in terms of registries and registrars there's a lot of tools available for DNS usage as well as DNSSEC things, and please, help yourself. We did them for the community, they are free and available… If anybody has any questions come see me.

João Damas: Okay, thanks for the reminder and I apologize for the omission.

Eberhard Lisse: I was going to say exactly the same thing. But my own opinion about this, as you all know I'm not an IT person by trade – I'm an, as you all know, a gynecologist. If a gynecologist can do it on a sickbed it cannot really be that difficult. It took tools, the tools are there.

João Damas: Now they are there.

Eberhard Lisse: I used plain BIND. Before Sydney we used plain BIND. Now you've got SPARTA tools, you've got OpenDNSSEC, there is really no excuse for any ccTLD to say "No, this is too difficult." Sorry, that's not the case. The tools are there, we all use BIND or something similar. The tools are there. It just takes a little bit of

time to sit down and read the documentation. It's really not that difficult.

João Damas: Yeah, the really big change when introduced DNSSEC is you go from a model where before you had the DNS zone, and if you published it it could stay there forever, it had no problem; whereas if you sign it, signatures have an expiration date. So the operation model, perhaps not for a TLD that has people dedicated to these but when you communicate this to your second levels, the operation model of DNS changes fundamentally in that you have to pay attention to DNS from now on. It cannot be something that you just leave there in the corner because it always works.

Eberhard Lisse: Alright. Are there any more questions? Okay, then the next speaker – we are a little bit ahead of time, which is not a bad thing – is Carsten Schiefner. He is now consulting with Ion DNS and not with DENIC so much anymore, and he is going to speak about IPv6 experiences from a registry's perspective. They run, Ion DNS or [Click Media] rather runs the backend for .cat so they have got some information.

Unfortunately they wrote this in a weird format – SVG. I tried to download two programs to give it in a PDF so Kristina could load it up. Unfortunately it didn't work so the remote participants will have to listen and be amazed by what Carsten has to tell us.

Carsten Schiefner: Thank you, Eberhard, and good morning everybody. Actually it was planned that Elmar Knipp, the Managing Director of Knipp Media.denic, behind ARIN DNS was meant to give the presentation. I'm actually just in quotation marks a "stand in." Quite a few know me with a different hat on and even worse, I only became familiar with the presentation this morning because yeah, it was…yeah. I simply haven't had any access earlier to this presentation, so if the flow of the presentation isn't that perfect my apologies please.

It's basically about practical remarks on IPv6 and DNSSEC from say a registrar's perspective and also from a DNS provider perspective, or from a DNS operations perspective. Content is first a short introduction to ARIN DNS; secondly some statistics from the IPv6 Day just the other week, and that happened just the other week; third, DNSSEC Made Easy; and yes the [FDS/PFN OC] which is, don't run away – not an ICANN acronym but this is basically the cliffhanger or the teaser for you to stay on, and I'm going to solve it later.

ARIN DNS is essentially a managed DNS service. Why would you use a managed DNS service? It's essentially about diversity. So really the core thing is if you don't have that already you may want to add some different management culture to your DNS service as a registry but also as a registrar. And you want to avoid software monoculture, because any kind of monoculture drastically increases the vulnerability of your service, because if your DNS

operations just runs on different machines but all the machines are the same hardware with the same operating system, with the same name server software, then if there is a failure or a buck in operating system or in the name server software then all of a sudden your entire DNS operations can potentially breakdown because there is a vector of attack at hand. So essentially the idea behind ARIN DNS is to offer a service where yeah, DNS operators can just add diversity to their own system.

Features are it's completely new written, it's written from scratch. It just does not use any say prior implementation or public software – it's really written from scratch. It's more than five years of development, it's close [chores] and it runs SAS service. So you just cannot buy or get the software; you can run the service or you can buy the service, actually purchase the service.

Obviously it has all the features that are currently being discussed and needed. It's unicast and anycast. It has v4 and v6 on all nodes. The nodes are currently ten nodes distributed across the planet in all continents, at those locations where from a network topology it's interesting to have a node. And obviously they're highly redundant.

DNSSEC is also built in and as it's positioned it's a premium service. So yeah, this is as I said a 24/7 managed service and it's also easy to include in existing name server infrastructures, and ARIN DNS and Knipp behind it is a neutral and independent service provider.

Coming to the statistics and our findings from the World IPv6 Day, I guess it was on the 8th of June. Why would you be interested as a registry in v6? Obviously I'm not going to repeat all the known facts already. The v4 space is exhausted. You are having users right now still using v4 and continuing to use v4. You are having users using a dual stick mechanism in their boxes, but in particular here in the Asia-Pacific area you will see an increasing number of users, registrants for example, that only can use v6 because the v4 space is simply exhausted. There are not any further IPv4 addresses available.

And so to just accommodate those needs, you really as a registry need to think about implementing IPv6 across all your systems, across the entire infrastructure you're running. And that is what I said – you need to really think about the infrastructure first and maybe then the applications come second, but you want to have a rock solid layer where yeah, your applications basically sit on and make use of the infrastructure.

These are some stats from the DE-CIX which is the largest internet exchange in the world right now, located in Frankfurt in Germany, and what you clearly can see is before the World IPv6 Day the traffic pattern was about .8 GBPS, and during the IPv6 Day it almost doubled really, which traffic is a generic traffic really. If you look at a certain top level domain, for example .cat, that is where we got some statistics – the red is really v4 and the little green is v6 traffic for .cat, and on the 8th of June there's no different traffic pattern compared to all the other days. And if you

look at the v6 traffic only that is even a clearer picture that there hasn't been any change, although the overall traffic during the IPv6 Day increased as I said by 2, by effect of 2.

So the conclusion essentially is v6 traffic in the DNS world is still much lower than in the application world, and maybe as a registry you just want to think about that at least 50% of your DNS servers should be v6 enabled. That just doesn't mean that they will see immediately a lot of traffic, however you would be able to accommodate upcoming needs by users of their service because they simply will switch over to v6 more and more. And if you will look at the IANA database, some TLDs still have only one IPv6 name server and many with less than 50%, so I guess that is one of the challenges of a TLD registry right now, to look after their own infrastructure and to come up with an idea of how to do IPv6.

DNSSEC Made Easy, that's the third part – better safe than sorry. Essentially in a registry you have two main core functions. One is provisioning your shared registry system which in the EPP world is our C-5910; and when it comes to DNSSEC you need obviously to sign as we've just heard from João already, you need to sign your zone. You could do that yourself or you can ask your name service provider to do that for you.

Another again example is .cat. The registry system, the SRS is core and the DNS is a name service provider. This is basically essentially the overall architecture, and we go to the lower left corner now which is essentially the hidden master name server

where the registry system actually provides the zone from [SRS2], and then via a [thatcher] it goes into the DNSSEC model of ARIN DNS where the entire zone will be signed, and later on being distributed in either the anycast cloud or the unicast service.

So in essence, if you do not intend to do that yourself maybe you just want to ask your name service provider to do that for you, which brings me to the last point of my presentation: the FDSPFNOC. What does it mean? It's not an ICANN acronym; it's the Free DNS Support Program for Non OACD Countries.

The program is one anycast cloud with five nodes. It's obviously, as it's ARIN DNS it's before NV6 enabled. It also could do DNSSEC for you if you like to, and the means of getting your zones into the cloud is either zone transfer or incremental zone transfer. The conditions are you need to be a ccTLD, you need to be not-for-profit; you need to be a ccTLD from a non-OACD country and you need to have less than 100,000 delegations. And the process is highly un-bureaucratic so either way, see me after the presentation or after these session actually; or see [Elmer] and I'm happy to talk to you.

And I guess that pretty much brings me to the end of the presentation. Thank you so much for your attention and maybe there are questions?

[Applause]

Carsten Schiefner: Oh, Nigel.

Nigel Roberts: Good morning, I hope you're as awake as I am or not. Interesting presentation, and diversity, genetic diversity in DNS is a very, very important thing – we've been looking at this for many years. One thing that puzzles me however is that this is closed source and software as a service. Without trying to impugn what you're doing, because I'm sure actually it's probably done, developed completely independently unlike some other DNS providers, but how do we know this?

If it's closed source and software as a service, how do we know that there aren't common vulnerabilities with BIND or NSD or one of the others? Because we can't see.

Eberhard Lisse: On what name server? Are you using BIND or NSD or did you adapt a piece of the software yourself?

Carsten Schiefner: Yeah, it's a completely written from scratch, completely newly written system from scratch.

Nigel Roberts: Then how do we know you're not doing the same thing?

Carsten Schiefner: I guess that is basically the catch 22 a little bit here. Trust me or not or trust us or not, but yeah – that is how the concept is essentially built. It's not open source so I mean obviously ARIN DNS or the team behind ARIN DNS, whenever there will be vulnerability obviously it's going to be fixed more or less immediately really. But third-party checking is simply not foreseen.

Eberhard Lisse: .cat is running on it.

Carsten Schiefner: Yep.

Eberhard Lisse: Since when?

Carsten Schiefner: As I've just learned about this presentation in general I can't really tell right now. I have no data.

Eberhard Lisse: But upon issue, right from the beginning if I'm not mistaken.

Carsten Schiefner: I guess, yeah.

Eberhard Lisse: So that's what I'm saying – it's not just something new. .na runs on it since two days, actually, but my reason was they offered us this thing free of charge. It increases my anycast provision to five anycast providers and so if something goes wrong it wouldn't really affect us that much. So it's a good thing.

On the other hand my experience with the host that they are is quite good, so that's why I thought I'll try it out. And running .cat is a top level and under ICANN conditions they have to have certain service levels since they have .cat. That's good enough for me as a small ccTLD to trust the system. Now I am not trying to make much advertising here, but as I said we like to offer, to make services that are free of charge; make them known so that small ccTLDs as yours and mine have opportunities.

Any other questions? None.

Carsten Schiefner: Okay, thank you. And again, I'm here at least until the lunch break and also during the lunch break, so if you're interested in the FDSPFMOC please see me and we'll set that up anytime soon, without any kind of application or bureaucracy or whatever. Thank you.

Nigel Roberts: Just before we go on to the next presentation and I hand it back to our learned Chairman, you'll probably note from the agenda that

this afternoon we have the return of the popular "Question Time." We'll have a number of registrars, well-known people from g space as well as cc space. What I'd like to ask is that those of you who are planning to come back this afternoon, if we can help get the proceedings underway a little bit by if you email some of the questions.

Now, the questions can be anything you like so long as it is vaguely technical- or policy-related; and if you want to give registrars a hard time on things like IPv6 or registry, ccTLD registry/registrar relationship from a technical or policy point of view, please send questions to me: nigel@roberts.gg. Or send them perhaps to Eberhard if you didn't get my email address, and hopefully we get two or three questions or even more to get the ball rolling this afternoon. Thank you.

Eberhard Lisse: Alright. Our last speaker before lunch is Jacques Latour. He's the CTO of CIRA. He is going to speak to us about their experience in IPv6. He also has the distinction of having the last presentation before lunch because CIRA graciously is sponsoring the lunch. The lunch will be at the Swiss café at the Swiss hotel so it's somewhere around that way – it's within two minutes of this. I'll just say it in advance so that he doesn't have to delve too much into that.

Jacques Latour: Okay, thank you. So I'd like to talk about our experience with IPv6. We started this around January timeframe this year, and I guess the scope here is doing IPv6 within Canada. It was quite an interesting journey because most of the ISPs in Canada didn't support v6. There was a lot of knowledge in Canada on IPv6 that wasn't there, so it was quite an interesting journey.

I just want to make one comment: there are a lot of people that talk about migrating to IPv6 or transitioning to v6. We need to coexist v4 and v6 together for a long time, so just a note. So when we did our v6 project, it was quite a big project plan that we had. The first thing we did is a lot of training and education on IPv6. The scope here is we wanted to participate in the IPv6 Day, the World IPv6 Day, and this presentation is mostly from an enterprise point of view – like getting infrastructure at CIRA and a portion of the registry to be IPv6 accessible.

So the first thing we did is a detailed assessment of our infrastructure. We had a lot of surprise there. We defined the objective; we didn't want to do everything v6 – we had specific components we wanted to do. We built a detailed project plan with lots of details in there. We did an architecture and design work and a lot of testing before going online, so that was about a six-month period to get done.

So the objective was IPv6 World Day, or World IPv6 Day, and the goal was to have our web presence IPv6 enabled; and also the IT Operation Team within my Operations Team to be IPv6 so that it could support infrastructure. We already had two DNS

secondaries that are IPv6 enabled, but the key thing is make the perimeter v6, make it permanent. It's not just a one-day thing so by June 8th I wanted the entire infrastructure, this infrastructure to be a permanent implementation. So by doing it permanently it means that we need to have the technology, the people, the process, the whole infrastructure to support v6. So that was the project.

We had a critical pact, which means that before you start anything you write a security policy to say "Hey, this is what IPv6 is, this is what all the security things you need to look at with IPv6," and that was probably one of our biggest challenges. Because there's not a lot of people in Canada who understood v6, even less IPv6 security, it was a lot of work to build our security policy. So we managed to do that and we made it public on our corporate website, so anybody can go there and download the template that we've written for IPv6.

In terms of transit, pretty much all the ISPs in Canada don't support IPv6 so that was a challenge, so we had to order new circuits to support that. And then we got the project underway. So in terms of our transit, one of our guidelines, one thing we did in the past is all the new circuits we're going to buy need to support v4 and v6. That was a surprise for us. One of our largest telco's in Canada, they managed to enable IPv6 on their existing v4 circuit but only for three days. So that was kind of hard to do testing and the whole program, so we did order new circuits with permanent v6 with a larger ISP.

So one of the policies we've put in place is we've asked all of our transit to support v6 and if they don't, when we go to renew we're going to cancel them. And we already told them that this is the way.

So like I said, the key thing was to build the security policy, but before that and then after that we had to do a detailed assessment and we discovered that a few of our load balancers didn't support IPv6 so we had to replace those in short order. We had to put those in the lab, test it; so doing a very detailed IPv6 assessment is really important but you have to find the right people to help you do that, so that was a challenge. So we had to replace some of our internet transit, and then when we discovered that some of the application we wanted to make available on v6 didn't work well over v6, so we had to fix some of that. But overall, after the assessment – that was pretty much at the end of January – we figured out we're in good shape to do World IPv6 Day.

And then the next thing we did is we started to do the architecture for v6 and then we wrote down the rules of engagement. So we've defined how we wanted to do this: keep v4 as is, don't touch it. The rule was dual-stack, so to support this we dual-stacked every component in the critical path to support v6. From an enterprise point of view no tunneling, so everything had to be native IPv6; so we're looking to corporate infrastructure, the registry – it's all stuff that we did need to v6 without tunneling.

We had one also, one IP guideline, and that meant we're going to use – because you have many IP addresses with v6, we're going to

use global addresses only. And then the ULA, that's a 64-bit part of the address, or the ULA's a temporary address, or whatever the (inaudible), so we use a global address. I think from a security point of view the hardest thing to do for us was getting over the fact that you don't NAT over IPv6, so that was quite a challenging thing, especially for security guys who did v4 for a long time. You NAT and they, I guess they assume that net is part of the v4 security. So after we got over that then the architecture design work worked a little bit better.

The privacy within the IPv6 addresses, so we've defined a few roles around that. We wanted to have temporary address; we didn't want to make the [Mackie] address of all the computers public over the internet, so we EUI-64, we didn't permit that. So we had some challenges around that in the beginning with the Macs internally – they didn't support that.

The other thing was to build an IP addressing plan. We found a lot of stuff around SurfNet and some other… Oh, I'm having a Microsoft moment here. So on the IP addressing plan we used RFC-3531, and basically if you haven't done your IPv6 plan yet for addressing you should take a good look at that. Use the left-most bit to define the segment and the right-most bit for the host. Once you look at the RFC, the first time you look at it you go "Wow, what the hell is this?" but after a while we really got to love that plan for v6.

In terms of IP address allocation, the HCPv6 is what we wanted to use internally within our infrastructure, and right now the Macs didn't support it but we did test last week the latest version of MacOS, and it does work with the HTPv6. It's got temporary address, it's got the whole thing, so that's a good thing within the enterprise to manage our IP addresses.

One other challenge we have right now is logging correlation with having temporary addresses. On Windows 7 the address can change every five minutes, it depends; but then when you do logging we've got to figure out how to relate that back to the exact user and all that. So privacy, this derives from we don't NAT. Because you don't want to NAT you need to have address privacy; because you have address privacy it's hard to track who's got what IP address at what time of day for logging. So that's something we're still looking at, trying to understand how we're going to manage all of that.

So more stuff around the (inaudible) mapping within the DNS. We had to learn all of the new routing protocol for v6 within infrastructure, like BGP, OSPF, HSRP, all of that stuff we use internally; and also net flow for data collection. And then the next step is we had to take, because v6 we didn't NAT, we had to take a lot of good care around managing the security to implement v6. By experience, when a keys in front of a firewall and they have a problem, they don't know what's going on they do "permit any any" to make sure the whole thing works. If you do that the problem with v6 is that you're not NAT-ed, you open up your

entire internal infrastructure to anybody outside, so that's one of the risks with not NAT-ing. So we need to have all the right procedures to manage this properly.

So we started by disabling IPv6 on all the devices within the infrastructure because some was enabled by default. We had to upgrade our IDS IPS with new device that did D packet inspection within IPv6. We enabled dual stack IPv6 on all the desktops and we had to make sure all the firewall and all that stuff actually supported that; and then for the logging, make sure that all the logs work properly.

So basically all the next slides in here talk about the security policy. I'm not going to go in detail with the stuff from now until the end, but basically these are all the things you need to look at – disabling to redo and all that stuff on the Windows machines. ICMP is totally different with v6 than v4, so there's a lot of work you need to do there to understand in the firewall what goes through, what doesn't go through and so on; and then fragments and all that stuff. So that was a huge learning curve for pretty much everybody to support the IPv6 infrastructure.

So this, all these slides, the stuff in here is available on our CIRA.ca, on our knowledge center. You can download the security policy. So we built a detailed lab, we got the whole thing working in the lab; a lot of testing and then once we were happy with that we did production. And basically it worked. So doing v6 is not just getting a bunch of people to go on the router switch's host and doing it. There's a technology, people, process aspect to

implementing v6, and on June 8th it was a success. Well, we had 480 hits over v6 during the day – not a lot of activity but it was working.

So I think maybe to answer one of your questions, IPv6 is a small project. It's a small piece of technology, it's fairly well understood and it took us five months to implement IPv6 properly. And DNSSEC is way more complicated, the project plan is way bigger – it's going to take us a long time to get all the components right to do IPv6. That's it.

Eberhard Lisse: Thank you very much. This was quite an interesting presentation as well.

[Applause]

Eberhard Lisse: Being a small registry, I don't care much about IPv6. My host has it we'll implement it; my host doesn't get it, we'll go out of business so we'll use your approach: "Either you provide IPv6 or we're not going to renew." That'll work, that'll work and eventually they have to do it anyway because IPv4 is going away. But if you want to do it it's good to have a plan and you can't just go and switch on your machine overnight. If it is a sophisticated operation where there is more than one Prime Minister involved, or

in a large place you need to make a proper plan. I think that's one of the ways of going about it. Any questions?

Gihan Dias: I'm Gihan Dias from .lk. I was wondering why you were so keen to say that while you're using IPv6 that you should not use NAT. Certainly I mean for servers I think that's true, but for like VCs and the client machine and stuff, because I do see for example the same thing you said – that having NAT gives you certainly some amount of protection from outside, by preventing people from connecting by mistake. So why not use v6 NAT for that type of requirement?

Jacques Latour: IPv6 was, I'm not the engineer that designed it but it was designed to be not NAT-ed. It's one host, one IP meaning that host internally and that host on the internet has got the same IP addresses.

Gihan Dias: IPv4 also was designed to have one host, one address, and that is some clue to which people did. But now that NAT is there and some people really can't live without NAT, I believe that people should be asked to use v6 NAT. And there's nothing wrong with that if you are.

Jacques Latour: Sure. Like I said, part of our learning curve was to understand IPv6. IPv6 is not IPv4. They're two different protocols. The concept, the design concept for each one is not the same, and part of the learning curve is we discovered that it's going to work way better if you don't NAT. You just have to get used to it. The firewall is there to provide security. In v4 if you don't program your firewall well you're opened up to a bunch of vulnerabilities. The NAT-ing in v4 doesn't protect internal networks because everybody gets [hacked in] so it's a perceived security mechanism – it's not a real security mechanism.

Carsten Schiefner: I agree completely with Jacques. In the end what happens is that NAT is creating today more problems than it's solving. You can see people using v4, widespread use of NAT – do you think the internet is more secure today than it was? I mean there was a time when the principal vectors of attack on hosts were port scanning. Today, you get viruses mainly via webpages and email and all this stuff, which goes through NATs as if nothing was there.

When you have NAT what you have is a lot of problems in applications that are not strictly server, client server, which have problems working. You have problems with new applications being generated because you force developers to introduce a large chunk of code just to deal with the problems of broken NATs, which are more common than the ones that supposedly work properly if they're there. So it has a lot more inconvenience than advantages if it has any advantages at all.

You might still want to have a look at firewalls, and I think what you're seeing now is that in consumer operating systems, the firewall is being put in as part of the operating system. But the firewall is definitely not the same thing as a NAT.

Eberhard Lisse: I would also think NAT is just giving us a false sense of security. It was used because there was not enough addresses and then we said "Hmm, maybe it protects outsiders to get in," but it doesn't. Much of your software opens ports from the inside that allows our outgoing, incoming stuff, so I would not insist on using a protocol that was developed or something that was developed because the address space is running out when we've got something which has an address space which is not running out, to do something for which it wasn't designed and then believe you're safe. It's complicated but you won't prevent a Conficker visit.

And I must say my own practice was affected. I've got NAT, I've got a firewall, I've got this, I've got that, and my secretary opened an attachment – whoops!

Khoudia Sy: My name is Khoudia Sy from Senegal in West Africa. I run an IPv6 test this year and our country participated in the World IPv6 Day last 8 June. I have two questions about your presentation. The first one is about the tunneling matters. I see you blocked all IPv6 tunneling matters – why do you only use native IPv6? And the second question is about security. I see in your presentation

you're talking about firewall rules. You say that we have to allow the "permit any any" and if we do so we'll implement security. It's native in IPv6 or not? Thank you.

Jacques Latour: I'll start with the second question. We don't want to permit any any, so I'm not sure. I said sometime when there's a problem that the keys, they go in and they go in and they go "permit any any," so that our first rule when you implement IPv6 is that you deny all the IPv6 traffic – so I think you may have misunderstood there.

The first question – tunneling. So when we did the architecture we said that internal users will now not be allowed to tunnel outside on the internet. They have to go native out, that was it. And then the goal was to provide 100% native dual stack infrastructure for that to work.

Eberhard Lisse: Any more questions? Okay. Thank you very much, quite interesting stuff and it's good to see that some ccTLDs are on the avant garde in that sense. Also, thank you very much for sponsoring lunch and that leads me to Nigel. Nigel is the character waving the lunch tickets in the back. On your way out, please collect a lunch ticket – we have 80 tickets, first come first served, and it's at the Swiss café at the Swiss hotel.

We should be back here by 2:00 – that gives us about two hours so we should be able to make it on time. Alright, thank you. You can applaud the lunch sponsor if you want.

[Applause]

[break]

Eberhard Lisse: Alright, then. So I hope we're all in a nice post-prandial stupor, the ones that arrived yet. The afternoon we're going to talk a little bit first about business continuity, and later on we'll have a little roundtable about communications after a disaster.

And Dave Baker, the CTO for .nz, I don't know which of the three companies you are owned by, will speak to us a little bit about their business continuity plan. They have had a few earthquakes these days, so not necessarily that they have been affected, but I think he will be able to offer some insight.

Dave Baker: Good afternoon everybody. Today I'm going to take our business continuity plan or BCP for short. Okay, what is a BCP? There's quite a long definition up there on the slide, but pretty simply it is how to recover your business after a catastrophic event.

So a catastrophic event could be an earthquake, a volcano, fire, flood, a financial problem, compromise of your IT systems and so on. In New Zealand, we have a major fault line as you can see from the diagram on the screen; it's basically running the whole length of the country.

And as it has been shown in the recent events in Christ Church, with the Christ Church earthquake, lots of buildings were damaged in the central business district, and a no go zone was declared over large parts of our CBD area. And that no go zone lasted several weeks, so no businesses or people could enter that area, except for the emergency services.

A lot of businesses located in the no go zone probably sustained little or no damage, but because a lot of them didn't have a BCP, they were severely impacted, because they couldn't get access to their key documents or IT equipment. So a lot of them since then are busily sort of putting in their BCP plans, and ensuring all their key IT systems have backups and are located at secondary locations.

Okay, so you need to implement a business continuity management framework or BCM for short before creating a BCP. The BCM framework contains a thorough business impact analysis of all mission critical activities and the services that underpin those activities, events, scenarios, risk analysis, and recovery strategies.

Implementing a BCM causes you to look strategically at your business operations and adjust where appropriate. For example,

we use to have our main data center located in Wellington, and our main office and IT support were also located in Wellington. After evaluating the risk of an earthquake in Wellington, we decided that that wasn't a good idea, so we moved our main data center up to Auckland which is several hundred kilometers to the North of Wellington.

So if an earthquake event happened in Wellington, then we would be impacted with our offices and staff, but the main IT systems should carry on functioning normally, because they're located up in Auckland.

Okay, inside our BCP folder, we have an emergency handbook which is just a small – small little booklet with useful information to be on hand in an emergency situation. We have a version of the BCP on a CD, a copy of the BCM manual, it includes the IT, the R-plan. Each member of staff and each director has a copy of the BCP folder. We have on handy in the office, one of our emergency kits located in Auckland and Wellington, and another one with our emergency outsourcing partners.

The BCP is essentially a guide with resources to enable you to create a plan to get your business up and running. In our plan, what we have – some of the key points identified with the symbols that you see on the slide, so an alarm bell means a risk to people, a light bulb would mean something to take notice of or consider, and a question mark, a policy or master instruction.

The business continuity plan is divided into two distinct phases; the emergency management phase and the business recovery phase. These phases link so that typically the emergency management phase leads onto the business recovery phase. However, procedures under the business recovery phase may be invoked in circumstances where an event has occurred that does not involve a present and immediate threat to people or major facilities. And I'll talk about each of these phases in the next few slides.

The emergency management phase of the business continuity plan is concerned with how to handle a crisis that typically involves immediate and present threats to life and/or major physical facilities, such as buildings or equipment. The emergency management phase is typically carried out in the first few hours after the event has happened.

On this slide, you can see a list of the key steps to consider, so it's not a must do all of these steps, but those are some of the main things that you should take notice of.

In our BCP we have documented who has responsibility for managing the emergency management phase and we also have a list of backups in case they're not available. After ensuring the safety of the people present, the first task in an emergency is to mobilize the team. This involves contacting the team members to appraise them of the situation and making arrangements to either assemble at a specific physical location or meet by electronic

means. And there's an alarm bell there which is sort of cautions you to travel carefully in a disaster situation.

In the event of a serious emergency, the first priority is protecting human life and welfare and then the minimization of the emergency situation, the elimination of the threats or of harmful factors, and the restoration of critical services. You need to continually assess the likely impact of the emergency event on the business throughout the emergency management phase.

Assessments should include consideration of the extent of the damage, who has been affected, and who remains available, and the likely time for recovery.

Okay, key staff should be kept appraised of the developments and how these may affect their own areas of responsibility. Obviously, if it's an event that needs emergency services, you need to contact them; we have details in the plan. We also have emergency evacuation procedures, but I guess you should really know those beforehand, rather than wait for an emergency to happen and try to find them.

Okay, so communications. Keeping people informed is one of the most important activities in an emergency situation. It is crucial that information is released in both accurately and timely fashion. Communication systems, telephones, mobiles are most likely to be down shortly after a major event.

For example in the Christ Church earthquake, the mobile phone systems were overloaded and were down. So you probably have to

look at alternate means of communication. And again in the Christ Church earthquake, the internet was largely unaffected and so that proved to be a useful means of communication through Skype, Facebook, Twitter, that sort of thing. And also that caused us to think a little bit as well, and so we went out and bought a satellite phone to have as a backup for our systems.

So at the end of the emergency management phase, you need to complete some reports. We have copies of these included in the emergency handbook, and every involved in the emergency should also keep a log of all the key events and decisions that they have been making. So they then should hand these over to the business continuity manager who will be taking charge of the business recovery in the next stage.

Okay, the business recovery phase of the business continuity plan is concerned with how to restore normal business operations. The business recovery phase can take place in the days following the event, and so this slide shows some of the key steps to follow, mobilizing the team, assess the damage, prepare a recovery plan, monitor progress, keep people informed and then transition back to normal operations.

So the first step then is to appoint a BCM manager. Again, this is detailed in the plan, and we also have a list of backups in case that person is unavailable or more than one person is unavailable. The BCM manager determines the status of the staff and mobilizes the team. The manager appoints all the team members and their decision, yes, it's final, it's their responsibility.

Depending on the situation and resources available, one person could have several roles in the same team or across teams. In our organization, we only five staff, so our five staff will be covering multiple roles.

Okay, in this slide and the next few ones, I've got a breakdown of the teams, and the key activities carried out by the teams. So if you're putting together your own plan, this could be completely different, have a different number of people, et cetera. So this is just the BCM team and basically responsible for overall coordination and decision making. The Facilities team are primarily concerned with damage assessments and recovery of the facilities.

Next up is the Information Technology team and obviously mainly focused with restoring of the computer systems and the applications.

And finally the Administration Support Team, so we're going to down one person and they're basically doing any of the tasks that's available, but other people could be seconded to the team to help out where required.

Okay, notification, this just depends on the event. For example, for our registry systems, if it's an operational issue, we'd normally have our first level of support being called out and then they would then escalate the call to either me or our chief executive, and if none of us are available, then it will go further up to the director of our company.

Implication of the plan, so that's normally done by the chief executive, but we have details in the plan of other people who are able to do that. We also have escalation plans in place and we also have a call down tree document, and we also have our main business phone number has a call down tree associated with that, so if the first number on the list of the company that manages our after-hours number can't get a hold of them, they'll go through a list of people until they find somebody to get a hold of.

Again, communications, so the media play an important role in disseminating the information to the public, however, great care must be taken in managing the context to the media to avoid the spread of misinformation and unfounded rumors. The BCM manager will be responsible for preparing press releases and regularly speaking with the media, regarding the organizations response to the crisis.

So training, testing and maintenance of your plan is important. Without these your plan is probably worthless. For our testing of the BCP plan, we conduct BCP tests – includes exercising our DR plans in the event of a Wellington earthquake. We've created detailed instructions on how our emergency backup IT supplier in Auckland can take control of our systems.

And what we found when we tested these was that we put together what we thought were the correct instructions, and then when they followed through those instructions, in some areas they did a few things in a slightly different order, which meant that they locked themselves out of key parts of the system, which sort of

demonstrated to us that although the instructions were clear to us when we put the plan together, you really need to test it with – if somebody else is going to be following those instructions, because they can easily misinterpret an instruction and do something that has a different consequence.

Okay, this slide is basically just detailing the steps that we went through in putting together our business continuity event scenarios, really important that you put these scenarios together and work out what your strategies are and contingency plans for an event.

And that's all, any questions?

Eberhard Lisse: Still post prandial stupor, but then an earthquake is obviously not exciting. Any questions? I mean in my country its simple, my servers are co-located with telecom. If they go down the prime minister can phone me anymore, so I don't worry.

But if you look at it from a serious aspect, I never thought about to actually run a scenario, turn that thing off and see what happens. Actually, pull the plug, kill the power supply, turn it off, see what happens, cold turkey.

No, we don't have earthquakes in Armenia but sometimes it rains. 2,500 names, three families depending on it – but when you start looking at it for the real business, I think one must have a plan. Any questions with regards to the plan from the registrars for example?

[background conversation]

Eberhard Lisse: Because you're there. Anyway, since we don't have questions, we will commence or proceed straight to our little roundtable. And I have invited Patricio from Chile, but he is prevented by natural disaster, flying through an ash cloud so he will come only tomorrow.

Then I had invited the two guys from Haiti, but very difficult to communicate with them, which we come to just now, they were just part of the reason why we're having this roundtable. So on short notice, we got Eleanor to volunteer to be the sacrificial lamb to sit there. Dave Archibald hasn't volunteered, but I'm going to volunteer him now, because he mentioned the word hurricane the other day to me. Nigel is going to moderate – no Nigel is not going to moderate, Steven is going to moderate, because he got flooded, or his place got flooded, and there was an earthquake and

I'm so much concerned about the business continuity plan in itself. And then there was Parkpoom Tripatana from – are you there? There you are there, sorry for not recognizing you, you also most of all come to sit on there, because you have been flooded quite significantly. Have you got a presentation or something? Not necessarily, but if you've got something you're more than welcome to start proceedings off, okay.

Now the reason what I wanted – what I wanted to actually do is you all know or have heard about that the (inaudible) has sent out a nice letter to Leslie, volunteering the information that he thinks that satellite phones are a good thing in such a disaster and we must all buy – or we should all – he suggests we should look into buying some.

Now I have spoken to Minmin from Dutch AP, who couldn't come last time because of the earthquake, who wanted to come this time, but he couldn't come, and basically I can paraphrase his email to me, what I should tell you. Two cables got cut, one right – one submarine right where the area, one terrestrial; and the DNS was not affected, from which he concluded that Japanese people were more interested in TV than in their internet, to be actually – to actually look what's happened.

What I would like to talk about is whether people was actually experience with disasters, actually think what type of communications are required after such a disaster. Patricio told me that he felt two way radios have been quite helpful because they – the people of some business – of the same office couldn't talk to each other.

The CTO from .nz, told me that some ISPs found satellite phone quite helpful when they could communicate and synchronize with each other to sort of bring up some equipment that cannibalize some parts and so on to get something going. And so I felt that I put this on the floor. But Michele had of course one question, now that I prodded him. No, you raised your hand at least.

Michele Neylon: You prodded me.

Eberhard Lisse: To the microphone and identify yourself.

Michele Neylon: Thanks Eberhard, it's Michele from Blacknight - a dirty filthy registrar, I'm fascinating by this discussion about business continuity, I'm also a little disturbed that by the kind of comments from some people about not ever, ever having tested us. Because from a registrar perspective, you know because we're expected to maintain high availability with respect to our services, so I would kind of naturally expect that the registries would be in the same boat.

Eberhard Lisse: Well we, .na does that, and we know of some of our compatriots that would have achieve to achieve continuity before they try to achieve business continuity – yes. He has in a face to face communication correspond – commented on that.

Anyway so without further ado, Stephen go ahead, make some fun.

Stephen Deerhake: Well, with regards to our experience in Samoa after the 8.1 earthquake we had about six meters of water come in pretty

quickly thereafter, this was early in the morning on the 29th of September, and our equipment that was located on island did not survive the seismic event. The ISP that we were located in did not have water issues, they just had seismic issues.

The second ISP on the island which also is the landline phone company had both seismic and water issues. The power company had major seismic issues, and really in the constellation of things that went wrong that morning, our losing our server by having it fly out of the rack and bounce across the floor really was a nit in the overall scheme of things.

The registry operation structured such that everything that was on island stopped working, but there was essentially rollover to the off-island facilities, but the secondary DNS servers and the registration system as well. So from an operational standpoint, it looked like nothing had happened external to the island. On island of course there was no phone service, no electricity, no water and no internet; so – and no mobile service either, so that was a different – that was a slightly different issue.

Insofar as communications going back to this business of sat phones, et cetera, et cetera, mostly what the island was using were two-way radios and it became a big federal – US federal disaster area over night; and once the Coast Guard got down there the next day, things started to stabilize a bit. And just today actually the high court has moved back into their courthouse building which was structure – had water damage, so that's one of the last ones, government agencies to move back into their headquarters.

I'm not sure how we want to proceed here, do we just want to work from right to left?

David Archbold: Oh, I can make some comments about different kind of disaster. Back in 2004, Cayman Islands, where I come from, had a Category 5 hurricane stalled right over the island for over 24 hours. And a Cat 5 hurricane is quite powerful.

Just out of interest talking about telecommunications, some of the landlines stayed active, one of two mobile networks lost it – actually throughout the height of the hurricane, although you were down, you had to resort to texting in the end, in the latter few hours. But both mobile networks were back up within 24 hours of the hurricane passing.

As far as satellite phones are concerned it depends to me on your set off. We actually – all our technical infrastructure is not island, and it's at various NATs in the US. So our problem was not keeping the technical infrastructure running, it was actioning, doing our normal management, and the connectivity between Cayman and the tech site, and that we managed by satellite phone, and we also, you can use the satellite phone with hit. It's a [beacon] connection as well, so we had both data and voice going over the satellite phone, which meant that we could maintain operations or resume operations pretty quickly after the hurricane had passed. So I'm in favor of satellite phones, but that's our particular situation.

Stephen Deerhake: Next up?

Nigel Roberts: Are you taking questions now or at the end?

Stephen Deerhake: Why won't we do them now? Eberhard, we need a mic.

Nigel Roberts: David, it's Nigel Roberts from Dutch EG. I'm interested to know what the costs of running data over that satellite link were for the period that you needed to –

David Archbold: I can't remember. The basic subscription for us is something, it runs about $100 per phone per month. That's what I – with the basic, and obviously then as soon as you start using it, your rates go up, but who cares in that sort of scenario, as long as you've got connectivity. So we pay about $100 per phone per month, I think.

Eleanor Bradley: Well, following earthquakes and hurricanes, I feel like a little bit of a fraud really, talking about the issues that we're facing in dot UK. But as you would expect for a large registry, we have fully – full business continuity and disaster recovery plans in place. And I would absolutely echo what Dave said in terms of the necessity to

rehearse those plans, not only because you find out the bits that don't work, but you know by rehearsing it, you can actually mitigate the problems in the first place.

And one of the first things that we as a registry or that we do routinely now is maintain a full risk register that not only considers the technical risks to the registry, but also legal and PR and all the other types of risks that we might face, and consider how those individual risks can be mitigated.

For us one of the things we've had to face, now the UK is notorious for grinding to a halt when it snows and in the last few years, we have had some you know for us, relatively heavy snow and has meant that high proportion of our work force haven't been able to get to the office. So while the technology is working perfectly, we're not able to service it, or service our customers. And so we've had to put in place, home working and other facilities to ensure that we can continue to maintain our service through the winter.

We also carry out desk based rehearsals and again find that extremely useful. It is very disruptive carrying out these kind of rehearsals, and so sometimes a desk based exercise can also tell you an awful lot about your plans.

I'm talking about the kind of issue that has been raised, where we've got some telecoms failure. The way we would handle it in the UK would very much to be to work with government agencies and use the established infrastructure that's already there. And I'm

very aware that that's very dependent then on the country that you're from. But in all likelihood for us if a significant failure or crisis were to occur, it would likely to be wider than the DNS itself. So for us it's very important to be recognized as part of the critical national infrastructure, which we are.

One of the things that – so within the UK, we would use an airway radio spectrum, which if the traditional coms network went down, which is something that already is a digital trunk to radio service, that the emergency services, police, ambulance, et cetera would use, and by hooking into the relevant networks in the UK, and having that role and being recognized as critical national infrastructure, we have access to that kind of facility.

The other thing that we've identified as being very useful is making contact at the local level so we are linked into at the local level the counter-terrorism people, because for example one of our elements of business continuity is that we have a generator that could keep us running in the event of a power failure. But if there was an incident that meant we weren't able to get fuel into that generator to keep it going, then that in itself would contribute to a crisis. So by linking into local networks again, they recognize the role that you play as a registry and ensure that you have the access that you need.

And I think just on the satellite phone recommendation itself, we would just that it's very narrow. It might be suitable for some registries, but not necessarily all, and it perhaps doesn't take into account the differences.

Stephen Deerhake: Questions.

Eberhard Lisse: Why are Europeans, Germans also always surprised when it snows? It snows. It snows every year. I know that there is fog in England. Why is it every year that everybody starts getting surprised of the fresh – wait a minute it's going to be winter, you must put winter tires on, why is that?

Eleanor Bradley: I don't know that we're surprised that it snows; we're just not very good at responding to it when it does.

Stephen Deerhake: I've got a question. How often or do you have a periodic schedule where you do a review of your plan, a formal top to bottom review, and I wanted to ask that of Dave as well.

Eleanor Bradley: Okay, we have a business continuity planning team that meet every two months, and we carry out – sorry, we have a business continuity planning team that meets every two months, and we carry out an annual schedule of fail over activities and business continuity activities to test everything, and then we would conduct a review – a full scale review on an annual basis, the but plan itself

is being tweaked all the time in response to new things that we're discovering.

Dave Baker: Yes, like for – we have a six-month testing of our plan with our IT team. We have a full review every 12 months and you know likewise whenever anything that's changed or amended throughout the year, that gets reflected in the plan as well.

Parkpoom Tripatana: For Thailand, our – there are some (inaudible) disasters in Korean Thailand for cities or the big cities in 2004, and last year we got a really big flat and about end of last year and then again at the start of these year at the same place. Most have like common situation like domain communication in Thailand is out, save for us of course, and then our vendor citizen happen; all the cell phone providers' signals are down and they cannot communicate.

Those are problem in Thailand and I think until now there is no plan to bring the communication up as fast as it should be possible because they have to wait until the service side provider brings the service side up.

One thing that I wanted to share is there is some research in Thailand, our professor, my boss actually do research about that after the tsunami in 2004, they do research about ad hoc network to help bring internet data communication up generally in their situation. The research called Dumbo. The Dumbo can – or

communicate between our east computer laptop and some – like some device and make a network temporarily and can connect until they file their endpoint, like satellite to communicate this outside the situation. Those are what we are doing.

For the registry purview, actually we never met like – this has never, never affect us, but we do some plan not officially, but we have a plan like we placed DNS outside our continent actually. We place at US, at our Euro because DNS is our heart of business. And I think most of us do that too, and our communication among our staff be like – be like we have plastic box that can connect to anyplace, any internet that are – that they can communicate with this other. That I want to share.

Eberhard Lisse: I want to know more about that?

Parkpoom Tripatana: About?

Eberhard Lisse: About the box, the magic box that – no, no, we don't necessarily do a one on one on this now, but this little box that can communicate with any internet, that sounds like a very cool idea, and I'm going to take about that.

Stephen Deerhake: Any other questions. Oh Nigel, surprisingly.

Nigel Roberts: I'd like to ask this question to each of the panel in turn, it's sort of a leading question I guess. But when you answer the question the way I expect you're going to answer it, perhaps you can just add your thoughts.

The question is basically this. Is no matter how big and how well developed your country is, let's say you're one of the most well-developed countries in the world, I mean we've seen this in Japan. Japan is a very well-developed country. Would you say it was unfortunate or risky to have all your DNS servers in the same country, in your own country, and any thoughts that you might add to that.

Stephen Deerhake: From my standpoint, the answer is clearly yes, it is a risk. Prior to the arrival of fiber, we were dependent on satellite links, and they're notorious for not working, working poorly and the dishes getting misaligned or getting blown off the island completely depending on the weather.

So yes, it's clearly not an acceptable solution from our standpoint to either have all the DNS on the island or all the registry operations on the island. It has to be split.

Eberhard Lisse: My own opinion is at the moment for semi-political reasons; I have no DNS in my country. I have removed intellectual property

totally out of the country until this is settled. But this is a calculated risk, and I hear what you're saying, but for me at the moment, I don't think – it's too big a risk so I can't take it.

But I think one should have any cast providers, two or three of them, and have them have a large proportion of the servers outside the country, and at least half of them, so that's what I'm thinking.

And I mean the bigger the IDDDD security and the less of an issue is because they have quite a plan, a cunning plan sometimes, but it's more for the smaller ones what do we do.

Eleanor Bradley: Do you want me to respond to Nigel's leading question? Yes, I mean obviously we would say that it would certainly be too much of a risk for us in the UK, but perhaps it does need to be country specific, and for us we mitigate the risk by having a number in the UK and almost an equal number globally. I mean it's got to be that way for us.

David Baker: Well, as you can probably imagine, we tend to follow what the UK does within our resources as best we can. So obviously I would agree with that. It's just that recently in both – in the second level in the registration level and then others, there's been a resurgence of an idea that named servers or facilities have to be only and entirely in the country.

And I think that one of the conclusions that comes out of meetings like this is that when somebody thinks this is a good idea, there are actually very good reasons why not to do them.

Stephen Deerhake: Any other questions, comments? Dave you got something to say there? Well, I want to thank the panelists, let me do that first.

Eberhard Lisse: Thank you very much, shorter than I hoped for but not a problem. What's the bottom line, do we need to buy cell phones or do we – no, what I intend to is do we need to get three quotation and send it to the (inaudible) so that they can pay for them or what?

Stephen Deerhake: Well, it might take as a bottom line, it varies by the situation, it's – if it completely is a custom, one off solutions for registries to think about. There is no one size fits all in this.

Eberhard Lisse: I think it's important that you have somebody outside to tend to look over as the primary, that doesn't require root changes, root zone changes. The root doesn't care really which one is the primary. All of them are authoritative, so that one of the secondary's can, oh, if you have a primary outside, it solve the problem; but that some of the secondary's can talk to the others that it pulls from them, so that continuity at least for the zone that

existed at the point in time, that there is some plan if there is loss of communication that some mechanism is in place that the systems that are outside take over and run it until communication has been reestablished.

Stephen Deerhake: That's essentially the structure that we have in place in Samoa.

Eberhard Lisse: Alright, you may step down, as they say. So the next one is Brent oh, and applause, of course.

[Applause]

Eberhard Lisse: Especially for Dave who got roped in on the last minute without warning, who was volunteered, yes.

Next one is Brent Lee who is standing in for Lester Kum from registry ASP. Where are you Brent. I saw him just now. Oh, he's hiding there. Have you got a presentation? Have you got it on a pdf? Oh, can you quickly email it out.

Brent Lee: Oh, okay.

Eberhard Lisse: Can you email it to Kristina.Nordstrom@ICANN.org, because she can unload it under the (inaudible); or you can put it on here, even better. Yes, we're not in a hurry, take your time, do one thing at a-

[background conversation]

Eberhard Lisse: Alright, we're still struggling a little bit with the technology, which is not unusual for Tech Day, so I think we can start. The resolution here is fine, just bring it up. No, it's fine, just bring it up.

[background conversation]

Brent Lee: Hello, I'm sorry for the technical difficulties as I'm running a virtual box on Mac. Microsoft doesn't, it runs very well in Mac. So my topic today is on the scaling needs and investment. As you know, a lot of ccTLDs, there are big and small ccTLDs. You have ccTLDs that are running 1000 names, 3000 names; and you have some that are running 100,000 names or millions.

So let's say you are a small ccTLD and you have about 3000 names. You have plans to grow your domains to about 100,000 names – what do you need? You probably will ask questions "How much do I need to invest?"

So one of the ways, it's missing in the slide; the slide is trying to say that infrastructure is actually quite expensive because over here you have so many components that you need to take into consideration. First of all you have your hardware costs, you have your application costs. You're going to buy applications to run on top of hardware; you're going to monitor it. You're going to have services running; you need to maintain the SLA. All those are costs.

And HA – high availability – everyone needs to ensure security, high availabilities on the server itself. Next you have redundancy and backups; lastly, monitoring. So all these components are costs and it could be various – it depends on your implementations. So if you don't control your implementation you could run out of costs very fast.

So actually do you need all these things? I think everyone needs all the securities, monitoring, etc. Other than that do you have any legacy systems that you need to integrate? Do you need that? For example, you're probably looking into integrating your registry system with domain company registrations in your [government], so do you really need that kind of service and how tight do you need to integrate? Whether you need to connect directly real time or is not real time acceptable? All those are part of the consideration where the cost is a concern.

Next is you probably ask "How much should I invest?" and thusly whether you can implement in stages. Usually all these are business positions. It depends very much on your implementation

and the needs from your management. I have an ideal setup of a registry here. You see the boxes here – those are very fundamental items for registry system. From the perspective of the end users, they can surf from the web or on mobile so they will come to your main site data center for certain applications management. When they come to your site probably they will have to go through some firewalls or (inaudible). If you have enough budget you probably will have IPS as well and you have IDS together with the IPS.

When it comes into your data center the components that usually you have are webservers, EPP servers, WHOIS, and you probably have some internal monitoring of all the traffic, etc., and you have your business intelligence tools to analyze the data. And you will have a database cluster with stand storage, fiber connectors and data backups, tape backups. You do daily backups, monthly, weekly kinds of backups.

Within your office you probably will have your help desk and you will have your own CRM service. Out of all this you probably will have multiple ISPs, multiple links to your data center, multiple sites out of the whole data center. You have a redundant site for DNS; most of you will do anycast. And you do multiple provider anycast – that's what most registries do. And some of them, they will have external monitoring to monitor whether the connections from the external source to the ISPs is working or whether there's any connection problem so that if some of the registrars complain you know how to handle the problems.

So this is very ideal for registries, but in actuality for a very small setup, you probably will cut half of the setup because let's say you have about 3000 names or 2000 names – you're running a private business, you're not government-funded so you have a very limited budget to run the TLDs. So what you usually will have is maybe you have a single ISP line or two ISP lines; you have a help desk, an open source help desk and you have web servers, EPP servers and some internal monitoring. You may not have some external monitoring. You will probably have multiple providers for anycast because a lot of the providers are giving free for certain numbers, so you are still pretty safe on the DNS parts.

But on the management part there is a bit of risk because the investment may not be there. So there's one other thing that a lot of people always think: "The more domain names that I have, I have to invest more money into the infrastructure." It's true partially but it's not fully true because if you look at all the systems, every system has a maximum capacity. So the most important infrastructure, it seems that that infrastructure can tailor to a certain amount of domain names. Only after a certain amount of domain names, then you will consider that your investment needs to be increased, so it's not a linear kind of relationship between the domain names and the platform costs.

So the thing that will escalate and increase your costs is actually on the factor of the risk – how much you want to invest to manage the risks. Whether you're a small ccTLD or a large ccTLD your risk is actually equal if your rules or policies are almost the same,

especially when you have a devious attack. The skill will be the same. Maybe they will start from a very small scale and then they will continue to upgrade the scale of attack, but the risk is always the same – they will attack you regardless of your size.

So to manage ccTLDs for small TLDs and big is actually almost the same and the cost is a concern. If you are private owned then it's going to be very tough for you to be a ccTLD administrator. So to manage the investment usually what you have to do is you have to collect a lot of data. The data collections give you ways to assess your risk, and you can present the data in a KPI format and set a level or threshold for your systems. A lot of open source or paid software allows you to do this kind of data mining where you can mine from different kinds of sources, whether it's your traffic data, your domain data. You can plot your data according to your domain names and your traffic to estimate what is your coming and growth.

Tools are available where you can collect multiple data sources, do all this kind of reporting. And tools are available for you to do intelligent alerts. Alerts, for example, you have some kind of performance alert. You can check out whether your growth and your investment and your traffic, whether it is linear or not; at which point whether the alert should trigger and tell you that "Oh, in a couple months you probably need to upgrade some hardware." That can be done – there's tools available.

This is a luxury for a lot of ccTLDs, especially the small ones. Dashboard is a very complicated thing if you want to implement,

but if you have the time and resources you can run this and it gives you a very good insight into your traffic, your data. For example you can check how many concurrent sessions are coming from your registrars; if you accredit more registrars, how many more concurrent sessions will be coming in. And you can monitor all those things in a dashboard format and you can know that whether the system can cope with those kinds of traffic.

So for investment purposes, the key things are you have to collect the data before you make the decisions on how much you want to invest or increase your infrastructure. And you identify the risk and you analyze which are the risks that you want to tackle first because some of the risks you probably can implement later. Then the last one is to upgrade based on the identified risks.

I've attended SROC. I find that SROC is quite useful for me to do this kind of planning. They do have a very structured way of managing, monitoring and managing the risk and the business continuity as well. So if you have time you probably can attend the SROC. If you don't know how to do all these things you probably can engage some consultants to expose some of the right tools. Some of the tools that you can use are Jawa – they have work tools. The [IRT], you can use that. And after you do all these analyses you probably will have to justify to your management; or if you're a government you probably have to take to the government "This is what to do," and if possible ask for some grants from the government to help you on the overall planning and investment.

That's all for my presentation.

Eberhard Lisse: Thank you very much, actually quite interesting. I never looked at it from this perspective. Just one or two remarks: I don't think you must look at it from the number of domain names you must increase, but from the number of transactions – whether if you've got 2000 or 10,000 it doesn't matter if you've got only so many transactions per hour because the storage requirements are enacted in the transactions.

But as far as external monitoring, .na is a small ccTLD as I said. We've got all (inaudible) plus we've got external monitoring. There is open source solutions. (inaudible) for example has got Landscape, there is Heartbeat. If my registry server doesn't talk to Landscape for more than five minutes I get an email. And then we also run some external, some monitoring admin or something that shows us the traffic.

So what I'm trying to say is you could actually say "External monitoring for a small ccTLD is mandatory" in your little slide. That is something that can be implemented without much resources. What is SROC?

Brent Lee: The secure industry operating framework that IPTLD arranged for us. One comment on the external monitoring: sometimes you need to monitor the rest points for each EPP comment coming in from

the registries and the registrars, so if you use external tools that may not be achievable because they don't track every, each EPP comment when they send and come back, what's the response.

Eberhard Lisse: This is not a requirement for na ccTLD because we don't have a service level agreement. So and most of the small ccTLDs don't really use EPP. We have got three external registrars I think – AP Mirror is one, [In Star] is another and three that uses the big ones, Markmonitor, they all use the web interface. And there is one inside the country that likes to play with EPP. Even if you have EPP you're not required to report to ICANN how long it takes so you don't have to monitor that. If the software (inaudible) tools doesn't do it yet but we are working on this to get that in. If the software can monitor it and report it I would monitor it just for interest's sake but it's not a requirement so this is not a problem.

Brent Lee: We do have experience that certain registrars sometimes, they complain that their EPP comment is slow. So that will help us to do some troubleshooting from external perspective.

Eberhard Lisse: Sure, but that's obviously a matter of bandwidth.

Brent Lee: That's an option.

Eberhard Lisse: As I said we have got, for small ccTLDs… For example our EPP service sits behind a big purpose. It's never been an issue, and in any case if you have got maybe ten transactions an hour, whether it takes a millisecond or a second it makes no difference.

Brent Lee: Well, the problem may be at your point. The problem may be at your ISP.

Eberhard Lisse: Yeah, the problem, the point is I don't care where the problem is if it's not an issue. If it takes a second that's not good as far as a service level agreement is concerned but it doesn't matter if you've got ten EPP transactions per day or per hour or per minute. If you've got ten per second then it becomes an issue, but you want to look at it. I would like to look at it but it's not an issue.

Brent Lee: Right.

Eberhard Lisse: But anyway, any questions? I don't want to monopolize this. Really, no questions? Alright, thank you very much.

Brent Lee: Thank you.

[Applause]

Eberhard Lisse: I found it quite an interesting presentation because it gives a bit of an idea of some things that we never think about when we do these things, and it's always good to learn from somebody who has done this systematically.

So now Nigel is on the spot and his roundtable, Questions to the Experts. Apparently two experts, no, one expert chickened out, the other one is coming.

Nigel Roberts: Rob Hall, come on down. Michele Neylon and Mikey O'Connor. Well, welcome to this part of the presentation which I hope is going to be a little bit more entertaining, keep you awake. It could go on for five minutes, it could go on for half an hour; this depends entirely on you the audience.

What we have here are three experts from gTLD land, amongst other things. I'll let them introduce themselves. There's at least one filthy dirty registrar. I said at least one but you were a self-described "filthy dirty registrar." I'm not sure the other people would like to take that appellation.

This is called Question Time. I've got a couple of topics here that we might want to talk about. If anybody would like to suggest just

a broad area that we could talk about then stick up your hand now. Anybody who's got any burning ambitions to find out what the aliens on the other side of the planet do? But EPP is something that we're very interested in. As a registry we're very, very interested in payment issues. Let's face it – you cannot afford to run your registry if you don't get payments in for the services you provide.

IPv6 is something that's at the forefront of what we're talking about. We heard a bit about that earlier on today, and so perhaps we can quiz the registrars on the panel about exactly what they expect from us in order that they can get on and do their job and we can help them. And perhaps the theme of this panel could be called "Perish the G-Men." So if I could just ask you from my left to right to introduce yourself please and say a little bit about who you are and what you do. Mikey.

Mikey O'Connor: That left to right stuff is really complicated. My name's Mikey O'Connor. I'm a registrant; I'm not a filthy dirty registrar. I've been involved in the GNSO policy process mostly at the bottom of the bottom-up process, so I've either been part of working groups or co-chairing working groups for quite some time. I don't have any clue at all about EPP, so I'll skip that. I don't care about IPv6 because nobody uses it yet, so if you can come up with some more topics I'll be fine, thanks.

Nigel Roberts: So you'll be the contrarian in this.

Mikey O'Conner: Nah.

Nigel Roberts: Michele?

Michele Neylon: Thanks, Nigel. I would be self-described as a dirty filthy registrar I suppose.

Nigel Roberts: Well, I've heard you use it about three times already in this meeting.

Michele Neylon: Well it's just this thing about registrars, we're always being picked on by various cohorts that we're always out to do no good. So I decided just to embrace that. Earlier today there was the vote on new TLDs, so personally I'm delighted that we've kind of got that out of the way – we can actually get back to talking about other stuff. And I'd be happy to talk about IPv6; not so happy talking about EPP in the technical side but in terms of people actually using it, more than happy to do so. In terms of the financials, again, that's something I'm more than happy to talk about. I'll let Rob talk.

Nigel Roberts: Rob?

Rob Hall: Hi, my name is Rob Hall. I am on the Executive Committee of the Registrars Stakeholder Group. I'm also their appointee to the Nomination Committee. We own the largest registrar group, one of the largest registrar groups in the world with over 108 registrars. We are the largest .ca, so I'm Canadian. We are the largest .ca registrar in Canada. And I have an interesting distinction I think, perhaps, that I'm also one of the founders of your ccNSO before it was called an NSO in that I was at the first ICANN meeting in Singapore. I've been to 36 of them and I'm not sure if that's a good or a bad thing.

But I represented .ca as the Chair of the Board of .ca for the first few years of ICANN so I also understand very clearly what a ccNSO is and how you operate and what the challenges are. I'd say just about 40% of our registrations overall are in the .ca and we have just over 1.2 million registrations. So Im happy to talk about policy, I'm happy to talk about technical – I'm a bit of a geek as well. So EPP, and I think the largest… If I can sum it up, the largest thing facing ccNSOs to a registrar, and certainly the larger registrars in the world is how do you stay relevant? And the harder you make it for us to deal with you the harder it's going to be.

Nigel Roberts: And it would seem that we have a chair squatter, so if you're sitting there, Eberhard, you are now on the panel. You have now been volunteered. Would you like to introduce yourself and what you do?

Eberhard Lisse: Well, for the 57th time my name is Eberhard Lisse. I have a day job – I'm a gynelcologist. And I have a night job – I'm an obstetrician. And I can boast to be the second oldest manager of a ccTLD after Oscar Moreno. I registered .na in 1991 and have been managing it continuously. That's the second longest without change in management, so the Deputy under Permanent Secretary for Paperclip son-in-law is trying to change this. But I don't think they will succeed.

I happen to be the Chair of the Technical Committee. We run a fully automated registry (inaudible) tools with the interface EPP. We have security with DNSSEC. What else do I do… Oh yes, and some people ask me how I manage to do all of this and I always say the day has 24 hours, not 25 hours.

Nigel Roberts: And I've just spotted a member of the panel who's been missing in action, so Graham, if you'd like to come on down.

Rob Hall: While he's doing that I forgot one piece of it that we can talk about, is we also own and operate pool.com which is the largest in

the secondary market or the leading domain space. And there's certain challenges around dealing with registries obviously when domains delete and there's a land rush for the deleting space, so I'm happy to talk about that as well if anyone wants.

Nigel Roberts: Okay. Before I ask Graham to introduce himself are there any other potential panel members I spoke to who are missing in action? I think the only one has actually left us and sent his apologies, but I'm just checking.

Graham Chynoweth: Graham Chynoweth with Dyn. We're a DNS provider relevant for these purposes, also a registrar; and we're intimately involved in the development of DNSSEC and worked with [Al] on getting .na assigned.

Nigel Roberts: Marvelous. So before I start picking on you, let's see… You've heard the various different expertise that the individual panel members have. Let's have a question, please, any question that comes to mind. Antoin? I'm coming up with the microphone, don't worry – that's why I'm still on my feet. That also means I can tower over them and intimidate them if I want to.

Antoin Verschuren: Antoin Verschuren, SIDN. For this panel I have a quite interesting question because this is the ccNSO Tech Day, so it's technical. And we have of course ICANN policy and specifically the GNSO ICANN Policy. One of the things that we currently are experiencing when we are talking about DNSSEC, for example, is that we see that on a technical level there is sometimes a need for a direct communication channel from DNS operators to a registry.

And in the ICANN model, especially the ICANN GNSO model, the registrars always say "No, there's no way DNS operators can talk directly to the registry because we say every communication must go through a registrar; and not only the administrative and the political and the financial communication but also the technical communication." And I think there is a latency in there that especially with DNSSEC might be a problem, and I wonder how the people on the panel think about that.

Nigel Roberts: That is an extremely good question to start with, and just to get this clear, by "DNS operator" you mean the technical guy who runs the name servers for a delegated domain name.

[background conversation]

Nigel Roberts: Yeah, so that's a yes. Now, from a historical perspective that's what the field technical contact used to mean. Now, since we read

the dirty filthy registrars have emerged out of the slime – sorry, I'm getting carried away with your metaphor, I beg your pardon – they've been putting themselves as technical contacts. But in some circumstances you go to your registrar and you put in your own name servers or the name servers of somebody like Dyn. So who would like to start this off? Graham.

Graham Chynoweth: Well, I think it's an interesting question and it's certainly something that we face. We're a registrar but most of our business has kind of historically come from the fact that we operate dyndns.com, which many people have used for a long time for dynamic DNS purposes. And so I'll disclose that my position at the company is not a technical one, it's a legal one, even though I'm involved in the technical issues.

But I mean it's certainly very pointed to us and something that we've continuously tried to figure out how to get out of the middle of because we are constantly confused on both ends of the spectrum, about both being contacted when we are the registrar but not the DNS operator and being contacted when we're the DNS operator but not the registrar. And I guess I'm not totally clear how the… I mean I'm interested to hear your comment on how you think DNSSEC complicates that, or cause I see it as a continuing challenge and one that I'm not sure if DNSSEC really changes fundamentally. So maybe I'm missing the point of the question.

Nigel Roberts: Okay, I'm just walking back over to give the mic back for a clarification.

Antoin Verschuren: Well specifically if you're the DNS operator and you want to get for example key material to the registry, and you have to go through the registrar there's a latency in the path. And sometimes you need to do this very quick. And whenever we are talking about models, for example, in the ITF or in other communities where we are trying to solve some problems of how to do this properly, we always get the feedback back "Oh, you know, the ICANN registrars are never going to accept that there is a direct channel, even if it's a very technical channel and if it's going over a VPNS secure channel or whatever. They always say 'It's communication from a registrant or DNS operator directly to the registry,' and by definition they will vote against it."

Nigel Roberts: Rob?

Rob Hall: Yes. So my simple answer is you're absolutely right – the last thing we want is our customers talking to many different registries, and I'll give you a couple examples of why. You have to remember that most of our registrants have a domain in more than one TLD, and I would say most of our registries that we deal with

do not offer anywhere near the same level of customer service – and I don't mean as us, I mean just differentially among yourselves as well. And so the last thing we want is to send a customer to a third party, being a registry, whose primary business is not offering customer service whereas ours is.

Frankly, registrars are in two businesses – one, we have to bill people; and two, we have to service them. I mean that's all we really do. So your biggest problem I've heard so far is latency and you know, being quite blunt, pick a registrar that doesn't do that then. I mean most registrars have real time systems, most serious registrars have real time systems that shouldn't be offering that; and if that is a differentiation point between registrars what you'll find is we'll use that to compete with each other.

So I can't imagine how you can send a command faster to a registry than through a registrar that's real time. If you're dealing with what we call a "basement Bob" in the industry, someone who comes home and doesn't have a real time interface, and reads their email once every day or two and sends your command to the registry – yeah, you probably should pick someone else. But the thing to keep in mind is we're the competitive level, we're the competitive layer if you will of the industry, and that means we should be trying to differentiate ourselves and offer different services and different levels of service.

But to be quite blunt, every time a customer of ours talks to a registry they come back at us, cause us more customer service time

and more grief than if we just dealt with the problem ourselves, because that's what we're set up to do.

Nigel Roberts: Thanks, Rob. I'm reminded of another industry that charges customers and then services them. Mikey.

Mikey O'Connor: Mikey O'Connor here. I know I'm supposed to be the contrarian but I'm going to back Rob up. I'm a customer of, I don't think I'm a customer of Rob's; I'm certainly a customer of Michele's and of several other registrars, and I'm a DNS operator. And the idea of actually having to talk directly to a registry about DNS issues makes my blood run cold, because I have domains in probably twenty TLDs, and the last thing I want to know is who those people are and how their systems work.

I want to go to the dashboard of my registrar, I want to use systems that are the same across all TLDs, and I will back Rob. I use Go Daddy. Go Daddy gets changes into the root within seconds, and that's good enough for me.

Nigel Roberts: Eberhard?

Eberhard Lisse: I don't think the problem is that you should go and look for a registrar that isn't a "basement Bob." The problem is that we have

got 252 ccTLDs, each of them makes their own rules, and a lot of them are basement Bobs. You don't for example run accredited for .na and we are well organized because we are a small market and it's probably not worthwhile. That's not the issue. The point is it's not that you must find a registrar who can do it; you must find a registrar who can do it and who wants to do it or who is accredited for the ccTLDs.

Now, of course basement ccTLDs won't do DNSSEC. In Africa there is only one, that's us, and then there is Mauritius that runs on the Afilias platform so that's not the issue. DNSSEC is not increasing the problem. The ccTLDs that do DNSSEC, they are organized in a way that I personally don't think registrars should monopolize this, but I understand their business model because it cuts into your margin because you have to put humans in to deal with the outflow. It's a business decision.

But as far as ccTLDs are concerned, I don't really care whether it's the client or whether it's the registrar. I prefer it to be the registrar but if the client tells me he doesn't want to do it what can I do? I can't force them and ICANN doesn't put me in a position to force them. I will go with what's least effort for me if I get the same results, and that would mean I would say "Okay, if the registrar does it it's fine." In the end we run (inaudible) tools, sooner or later (inaudible) tools will have this included and then I'll go with the flow – whatever the server does that's the way we go.

But the point is you've got 252 ccTLDs, each of them have its own rules. Some competent and big registrars don't accredit in the

smaller ones because it's not worth the effort; it's more cost than it actually generates because you go for the mass market and then you are not a viable alternative.

Rob Hall: Can I take issue with one thing you said? I don't think it's that we go for the mass market. So I'd be very happy to connect all 250 if there was one damn interface and one damn bank account I could do it through, because sending three different characters or two different characters after the dot doesn't matter to me. But you're quite right – when you make it difficult and unique to deal with all the 252 of you, we do a cost/benefit analysis: how many are we going to sell?

But I think it's one of the challenges you have, because the nice thing about the ICANN domains is it's the same conduit I have, I get access to all of them; for the most part it's the same interface, and as you see new gTLDs roll out that's what you'll find the Afilias' and the NeuStar's pitching to their clients – "We already have all the registrars signed up, come to us." Rather than sending us .biz or .info they can send to us .pizza, and if it's that simple we're going to sell as many as we can.

So I think one of your challenges is to make it that simple. And I know you're sovereign and I know you want to be deliberately different, but the more you cannot be on a technical level the more registrars will sell your product.

Antoin Verschuren: So again, I really appreciate, I really understand your problem that every registry is having their own systems and they are having their own administrative rules. And that's why I understand what registrars really want. What I'm talking about is really the uniform interface; for example, if you want to get key materials to the parent, why not use DNS for example? Nobody is ever going to interfere with the protocol, but we're not allowed to create such a protocol extension because it doesn't fit the ICANN rules.

Michele Neylon: Hold on a second – what have the ICANN rules got to do with ccTLDs? I'm sorry but both yourself and Eberhard keep on going on about ICANN rules. They're completely irrelevant. Unless you've signed some kind of bilateral agreement with ICANN they've got nothing to do with it.

Now, a registry and its relations with the registrars and the registrants – that's a totally different conversation, because every time my local ccTLD operator talks to one of my registrants, it creates about five to six hours of customer service headaches for me. But that's a totally different conversation.

Antoin Verschuren: Exactly. I'm not talking about talking; I'm talking about designing a protocol where things from child to parent go up automatically, every interface the same.

Rob Hall: That's exactly the problem. So I disagree with my fellow colleague Michele here. ICANN has everything to do with what you're doing. So while he's technically correct that they don't control what you do the market does, and in the market I'm going to take the path of least resistance to make the most money. I mean frankly that's, being perfectly blunt that's what we do.

So of the 500 TLDs coming, or the 1000 TLDs coming, if 300 of them make it easy those are the ones we're going to sell. And if someone wants to go out and create- I'll give you the perfect example: .xxx. They've extended EPP in their own unique way because they need to prove you're a member of the adult community, and that causes us heartache. Now I've got to go and customize an interface to them even though it's Afilias' backend. I'm having to create a separate EPP connection on a separate EPP protocol with a separate bank account, and a lot of registrars are saying "I'm not going to do it."

So I think unfortunately the market forces apply to you as well. So while you don't have to use standard protocols and you don't have to use standard policies, the larger that body of work gets on ICANN and the more domains are there, the more you may want to go that route.

Now, I think you've got a unique opportunity over the next year or so to change that. So I think if you can come up with a better mousetrap… I remember sitting down with the .xxx guys and saying "What are you doing with your protocol? How are you going to manage this and what information do you need?" and

wishing there was a standard, because you know, we handled the auctions for .asia in Pool, we did .co, now we're doing .xxx. The reality is they all need roughly the same information but they all come up with their own way of getting it, and it drives us bananas.

You have a unique opportunity in that most of the time you're collecting information that is probably greater than a typical generic like .com. If there was a way to standardize that, even down to the field name, that would help us greatly and you'll find more of us wanting to deal with you. But as long as you're unique then we get into a cost/benefit analysis and you're beat.

Torbjörn Carlsson: Okay, it's on the same topic. I'm Torbjörn Carlsson from .se. What we have seen so far because we have launched DNSSEC a couple of years ago, and the only way to send key material into the registry is through the registrar via EPP of course. And the problem is, and so far I think we have been on the topic but sometimes the panel I think misses the issue – the issue is the key material, how to send the key material into the zone file.

And in our registry/registrar agreement, so this is also a legal topic in this. Our registry/registrar agreement says that a registrar should use EPP, there is no other way, and so far we have about 25 registrars supporting DNSSEC. It means that they can send key material to the registry. Most of those registrars, in total we have 150 registrars but those 20, 25 who can handle DNSSEC, they are also operators, DNS operators, and that is the first problem: all of

them only offer the registrant that they must use their DNS service to handle the key material.

So I think what we will do when we will change the registry/registrar agreement is that we will force all our registrars if they're going to offer DNSSEC, they must also offer the registrant to handle another DNS operator so they can handle the key material for the registrant even though the registrants choose to use another DNS operator than the registrar. Are you following me? That's the first problem to do.

The other thing is more technical. I think all registries together must find some sort of technical solution for commercial DNS operators who are not registrars, or for small, for private persons who are running their own DNS. What we're looking into is setting up some sort of web interface for those who are running their own DNS service so they can send the key material directly into the registry.

For those who are going to handle hundreds, thousands of zones, having that then we must have some sort of API. We can use an open source but that could be some sort of "EPP Lite." But I don't think that's a good idea if the registry uses their own solution to solve this problem. I think we should go through together and do it in some sort of standard way.

So the more registries who handle this the same way, the better for the registrars and of course the registrant. Okay, that was a long question.

Nigel Roberts: Thank you. Mikey, you were next I believe.

Mikey O'Connor: This gets back to my blood running cold as a customer. I'm an end customer, I'm a registrant, and the concept that I would have to deal with different providers for each set of keys for DNSSEC makes me crazy. I really don't want to do that. I'm fine dealing with a couple, but if I'm in ten or 15 or 20 or 100 ccTLDs and I have to use a different process and different software to get to each one, that's crazy. Is that what you're proposing?

[background conversation]

Mikey O'Connor: Well, so what are you proposing if not to put an intermediary in there so that I don't have to do that.

Nigel Roberts: Okay, Roy.

Roy Arends: Hi, my name is Roy Arends; I work for Nominet. I'm just here to explain the issue and I'm not for or against it. The idea is that a DNS operator can have some automated tools, just for instance like you have SSH that generates keys automatically when they are not

there; so can for instance a DNS server or the operator for the DNS server instruct the DNS server to generate the keys when they are not there. They need to be rolled over; when they need to be rolled over there needs to be a message going somewhere – this can all be automated.

And the idea is that an automated message goes directly to the registry. It doesn't need any configuration, it can all be automated. This can all be discovered if that makes any sense.

Mikey O'Connor: Let me just be continually stupid, because I don't- You want me, the registrant, to do that? Think about the-

Eberhard Lisse: Please, everybody use the microphone because it's not going to go over the remote unless we speak into the microphone.

Nigel Roberts: Perhaps I can just- I don't think the issue is the registrant has to do this. The issue is that the DNS operator, the person running the name servers may not be the registrant. Now when that is the registrant, or the registrant has chosen to use a third party-

Michele Neylon: Nigel, Nigel – this comes back to the basic problem I have with DNSSEC. The people who designed it completely ignored operational reality and it is fundamentally flawed. I'm sick of

hearing registry operators and DNSSEC fan boys going on about what they're expecting registrars and registrants to do to make an imperfect, flawed system work for them. Why on Earth didn't they actually think about this when they were designing the damn thing?

Nigel Roberts: Okay, so Michele thinks DNSSEC is flawed. Who doesn't think DNSSEC is flawed and would like to speak about that? Roy?

Roy Arends: Hi Michele. Roy Arends again, Nominet. I happen to have something to do with the DNSSEC system in the past. I agree that the DNSSEC doesn't win any beauty contests and I disagree that it's fundamentally flawed. The model of DNSSEC follows completely the model of DNS; it's a hierarchical system. DNS was before there were any registrars; DNS was there before there were any registries, so we're dealing with a very old system and we had to back fit DNSSEC on top of that. Okay, I'm getting overruled by Warren here.

So I disagree that DNSSEC is fundamentally flawed. I think it's a very clumsy model. We have registrars… I'm not saying that registrars are clumsy; I'm not saying that this model itself is clumsy. The problem is that the registry/registrar model doesn't follow the DNS model exactly, so hence a bunch of us were thinking to have an automated system amongst various registries,

including gTLD registries – to have an automated system to enable the influx of DS records into their registry system.

That can be through registrars or direct, whatever's possible. I understand that an end user might not want to do that. I can also imagine that an end user might want to have an automated system in place. So I'm just saying that a thousand flowers bloom. I know you don't want registrants to talk to registries – I completely understand that model. But this is just a proposal. It's technically viable. This has actually nothing to do with DNSSEC being flawed; it's just a method to make the process more automated and a little bit more easy. That's it.

Nigel Roberts: Okay, so Roy says it's not flawed – it's just an awkward fit on an old system.

Russ Mundy: Russ Mundy from Sparta, one of the DNSSEC bigots around, I guess. I'd like to also point out that not only was DNS created before the registry and registrar model; DNSSEC was created before the registry and registrar model. And so if one wants to get into a finger pointing contest, which I really don't, but you know, why didn't other folks be thinking as we went along?

But I have to say when the initiatives were really being pushed forward to say "Okay, let's get DNSSEC out there," there was a huge gaping hole recognized with respect to the registrars that

things didn't fit and work as hoped. And now there are a number of different relationships that don't fit real smoothly, okay, in terms of who's operating name servers and what they're doing.

Now one of the things I would also like to contribute here is that the operation of a name server for a zone is not inherently linked to either the, if you will, "owner" of the zone – the registrant – or the registrar. It is a separate functionality and people need to think about it even though the dominant use is in a particular way. It's a separable kind of function, and if you think about it like that you've got to get names into a name server that fit what the desires of the holder of that name wants. That's what we have to do as a community and work together to do it.

Nigel Roberts: Okay, so Warren said he was just going to bitch so we'll take that as read. Let's wrap up DNSSEC and see if we can do a couple more topics. Eberhard?

Eberhard Lisse: I'm wondering who's addressing what audience here? We are ccTLDs, 252 of them – how many of them are doing DNSSEC? Twenty? 19 of them are big, they are non-problems. The one that is small, we are no problem. The problems are the other 49 African registries of which 40 you don't reach or 20 or 10. You reach them but your money doesn't reach them or your invoice doesn't reach them, or your money reaches them but you don't get

a receipt. To talk about DNSSEC in the ccTLD forum is not the right target.

Rob is quite right – his business model is different than our business model. You first have to service the local community. Most of us do not run like .co or some others are fully commercializing when there is a big incentive, a commercial incentive to lay back and enjoy what Rob has in mind for them. Smaller ones like us, I would love to get lots of registrations from Rob but I can't afford the investment it takes to do that. I cannot afford the lowering of the price that would do that. I have nothing against the model but I don't think we're the right audience.

Of the 252 TLDs, which ones are doing DNSSEC? The big ones – they are not a problem. I think you can register if you want with .de or .se as easy as it comes. You make a plan with them, you sort it out; there are enough registrations coming, you will even put human resources into it. But once it's developed it's done and it doesn't need any human intervention. They are not the issue – the smaller ones – and we will not be able to solve it by saying we must have protocols.

We have no contract. ICANN cannot tell me one little bit. Even if my application to join the ccNSO is approved it cannot tell me what to do unless it's a PDP policy and then I can still resign from the ccNSO and I'm still not bound to it. And that's the attitude many of us may volunteer. Half of us don't know any better; half of us don't even exist, they're just on paper so that somebody can run it. It's 252 different models and that's I think the problem.

Nigel Roberts: Eberhard, thank you for the summary there. I think we'll leave DNSSEC at that; I think we could go on all day and I'd like to take a second question from the audience.

Eberhard Lisse: We've got seven minutes left and then DNSSEC for Everybody comes up.

Nigel Roberts: Warren, I'm sure you've got something that you've got on your mind.

[background conversation]

Nigel Roberts: Okay. Mikey, in that case, if you've only got two minutes we'll stick with what we were, or…?

Mikey O'Connor: Well, one of the nice things of living in my own little world is that I can make up questions that people haven't asked. A couple of points came up that I thought I'd amplify a little bit. The phrase "deliberately different" got thrown into the conversation earlier in terms of the technical systems. I think "deliberately different" in the future is going to be a problem, both from a business

perspective in terms of talking to an intermediary who wants similar; but also one of the things that I'm doing is co-chairing the DSSA Working Group, the cross-constituency working group about security and so on. And heterogeneous systems, different systems like this present a bunch of really interesting and difficult security issues that you all need to think about.

Another point that I think you know but it doesn't hurt to amplify is that the registrars represent a distribution channel and if I were in your shoes as business people I would want to be making it easy rather than difficult to use that channel. And then the final point, no, that's it. The last point is my final point.

Nigel Roberts: Alright. What I was going to do, I was going to go down the panel and ask for some final comments before we do our closing.

Mikey O'Connor: Those are mine.

Nigel Roberts: And I'm going to take those as yours. So I'll go to Rob next, because he's got a point that seemed to follow on from that I think.

Rob Hall: Yeah, I want to comment on a comment that the gentleman at the back of the room made earlier, which was about contracting and

basically putting things in the contract that forces a registrar to be the same.

I think you struggle constantly, and from what I've seen, country codes struggle constantly as do some gTLD registries, but more so the country codes, with this idea that all registrars should be equal. And I think what you'll find from most registrars is we want to all be treated equally but please don't try and make us all equal, because we're not. If there's demand for a service from my customers I'd like to innovate and satisfy the demand, so in the example back there of the gentleman, that 25 registrars and not one of them was doing real time DNSSEC – hey, that sounds like opportunity to me for a registrar to go in and do that and the other 25 can do it or die.

But remember, we are the competitive level and that means some of us need to win and some of us should lose, and some of us should innovate things that you haven't thought of. But that's by design, and I know it makes a lot of people uncomfortable and perhaps I'm being more blunt than people are used to but that's what this level of registrar in competition is. And it works – the market will sort itself out, but you've got to be gutsy enough to let it be free and do that.

So we're enablers of innovation if you will. Please don't try and stifle it by making us all the same, either contractually, policy, or technically.

Nigel Roberts: Thanks, Rob.

Graham Chynoweth: Yeah, I guess just a couple comments. One is certainly that's something, that type of market demand is something that Dyn as a company is focusing on. I mean part of our roadmap is to have a fully automated mechanism to make sure that the keys are passed quickly, so I think that's something we're trying to do. It's on our roadmap to get done; it's not done yet, but and I think kind of tied to that is kind of market awareness. And this goes to the comment that was kind of up close here, which is there really isn't enough awareness about the differentiation between, certainly in the consumer's mind, between what a registrar is, what a registry is and what a DNS operator is.

It's completely opaque in most people's minds, and the only kind of hopeful commentary I would have on that is I think as we roll out new gTLDs the market for DNS services is probably going to expand and there may be a greater appreciation for a DNS operator as a DNS operator, independent from those other functions; especially as we get kind of registry services providers backing up companies which are registries but which don't do any technical work themselves.

So I think that as we move forward there'll be more transparency on that topic rather than less. And I can say that the market forces are at least driving this DNS company which also happens to be registrar to offer those more specialized services which might

make us more attractive to folks who want to focus on a particular ccTLD.

Nigel Roberts: That's very useful. Michele?

Michele Neylon: Thanks, Nigel. I'd like to go back again to that point, and I'm not sure which registry that gentleman was from but changing a registrar accreditation agreement to force registrars to do something – that is something that always upsets me. Registries and registrars should work together, not against each other. While I realize that not all registries like the idea of collaborating with registrars the more open minded ones and the ones that I actually like to hang out with kind of do.

Encouraging registrars to offer services is fine; forcing us is a bad idea, whether it's a ccTLD or a gTLD. I mean Rob and I may disagree at times on certain aspects of it, but from my perspective with new TLDs coming down the line, every single tweak in EPP, every single minor change in policy is going to make various extensions less attractive to me. And it's the same thing with the ccTLDs.

I mean some ccTLD operators are fantastic to work with and don't cause me headaches; others cause me such headaches that if I'm away in a different time zone I have to be frantically following up with various people to try and make sure that that ccTLD operator

actually does make sure that domain resolves even though they've been paid three times. Thanks.

Nigel Roberts: Thank you. Do you want to say anything, Eberhard, or do you want to take the Chair?

[background conversation]

Nigel Roberts: Okay, so in that case you'd like to come up and do the closing remarks? Thank you very much, and thank you for listening to Question Time again.

[Applause]

[background conversation]

Carsten Schiefner: Okay, so Eberhard always asks someone else to close the meeting so I was picked out today. I was thinking in the beginning a little history of this Working Group, and it's fantastic because it was quite a long time we started this approach and we are still meeting in a full meeting room of people. So it means there are still some

topics we can discuss. And maybe the importance of this Working Group could rise at a time because as you know, today it was agreed that the new gTLDs will appear and it's highly possible that those guys, those new folks will face the same problems as we faced or we are facing now.

So there might be some sort of direction we can go where we can a little bit expand beyond the borders just of ccNSO and maybe we could make the meeting a little bit broader, because I am pretty sure that we will have a common or shared problems with those new guys.

And anyway, than you very much, all, for participating. I would like to thank all the presenters for their excellent contribution and of course I have to thank our distinguished Chairman for organizing this great event again. Thank you very much. And of course, I cannot forget the Secretariat Kristina for helping us with that. Thank you very much.

[Applause]

[End of Transcript]