DURBAN – Tech Day
Monday, July 12, 2013 – 11:00 to 17:00
ICANN – Durban, South Africa

SPEAKER: This will be the Tech Day session for ICANN 47. The date is Monday, July 15th. The session will run from 11:00 am until 5:00 pm.

EBERHARD LISSE: Good morning everybody. If you could take your seats please? Nigel, could you take your seat too? Welcome to Durban, welcome to yet another addition of Tech Day. For those who don't know me I am Eberhard Lisse, I am the Manager of .na and I am the Chair of the ccNSO Technical Working Group which organizes the Tech Day on every Monday of the ICANN meetings.

I just want to say a few words about our Agenda and a few housekeeping things. Is Tom Barrett in the room? So our first presenter has not managed to arrive yet. Not a problem, we can just shift the presentations a little bit. Our second presenter has also not arrived yet, but… Well, we're in Africa.

In the morning we have a more or less African focus. Tom Barrett is not really on an African focus; he will speak out about the Trademark Clearinghouse. Pierre Dandjinou is the VP for African Affairs of ICANN. He wants to give us feedback on the African DNS Forum that was happening for two days before the ICANN meeting.

I'm going to talk a little bit about fun and games with CoCCATools. My colleague, Ben Fuller, from .na will talk a little bit about how to care and groom emerging registrars. Then Nishal Goberhan – is he here? – from AfriNIC is supposed to talk about their experience setting up Internet exchanges and the L-root initiative. And [inaudible 01:43:22] that they have done over the last few months.

And then this time we have managed to get ourselves a lunch sponsor again so we will have boxed lunch here and they will give us a little presentation. They are a data [sister? 01:43:07], which is a data center. So since none of us are their clients I felt not really bothered by asking them to give us a little presentation before it.

In the afternoon Sunday A. Folayan from .ng will talk a little bit about ENUM. Rod Rasmussen, who is a Member or Chair of the Expert Working Group on Directory Services, the entity that is looking at replacing WHOIS, will speak about that and also about botnet mitigation. Then Martin van Horenbeeck from Google is going to speak a little bit about DNS hijacking. We have the usual host presentation by Theo Kramer.

Chris Hesselman from .nl is going to talk a little bit about a DNS set of tools that a colleague of his developed at .nl, who isn't here. So he will speak about it in his place. Nigel Roberts will speak a little bit about bitsquatting, which is a variant or type of squatting that has recently come to our attention, which is very interesting.

Then we will hear a little bit about certificate authorities, especially in the light of recent developments oversees in America with regards to data capture by governments. That might be interesting. And then my personal highlight: Richard Lamb will give us a second updated version on the manufacturing of HSM in a toaster oven again, which was the coolest presentation we had in Prague and for quite a while since then.

And then as usual – though somebody else this time – Andre Filip, Deputy of the Working Group, will make closing remarks. So if Tom Barrett is here now… Some people have arrived. He isn't? Pierre Dandjinou is also not here so it will probably fall upon me to start the proceedings. Let's just hang on for a second. I'll quickly find my presentation.

We recently had occasion to interact within the Namibian Competition Commission as a frivolous complaint was laid against our company with regards to abuse of a dominant position. As it turns out, our turnover is so minimalist, in 10% of the threshold that we don't even fall under this regulation. However, we made a presentation to the Competition Commission to show then what we're doing and so they understand better what they're dealing with.

And while doing that we noticed that we have a researcher who moves between registrars. And I was wondering, can we graphically display domain transfers? I will then show a little bit about transaction patterns. I'll come to this when I reach the topic and then also registration patterns, which is basically how many we have got in the second-level domains and things like this. And then we have a few minutes for discussion.

So that's a very good picture but what does it mean? It shows every transfer of a domain name since the 1st of January 2008, when we really started… We became live in December 2007 but for purposes we actual traffic renewal registrations happened on the 1st of January. We only moved in the middle of December but there wasn't really any action going on with one or two registrations. The other [three? 00:20:08] migration of the legacy system.

This basically shows every single registrar that has had a domain transferred to them, in black. Unless it is also losing registrants and then it has a different color. The arrows are the same color. Of course this picture doesn't really show you much unless, if you look on the bottom right, you'll see a cluster of domain names, which are the international registrars.

We have done this and I wanted to go into a little bit of detail to show how easy it is, even for a simply gynecologist like me – not a programmer who has a degree or knows programming languages – to look at the data that we have got and display it in a way to analyze it or even make predictions. I remember very well on our very first Tech Day in Sao Paulo, many years back, Jay Daley now from .nz said: "We have the data, we have the resources and the know-how, but we don't really make use of it."

At a later meeting in Cairo we had a presentation from .cl where it was attempted to make economic predictions. In other words this data mining. With 3,000 domain names like us it doesn't really give us much information as far as predictions are concerned, but if you've got 500,000, 800,000 or 15 million, of course you can use this data.

These are open-source tools that I use: Graphviz, Perl is the language, some modules that you need to interface with the database in GraphViz. Since our database is on a computer that we have access to over a network we use SSH Tunnel and on the Mac you've got a graphical interface to set this up so that the transfer is secure.

I use LaTeX to write my things and I wrote also with the Beamer module to write this presentation and I think it looks quite good. During this I figured out what the difference is between vector graphics and raster graphics. Raster graphics is little pixels where especially if you have got round or ovals, if you go to high resolution the picture becomes ugly. Vector graphics is much better; the outcome is much high quality.

I will show you. And you can do this with a program called Inkscape, which is reasonably quick, especially if you call it more than once. But I also found this library, LibRsvg, which can do this very nicely and just Google it and download it; it works on almost every reasonable computer system.

So let's go a little bit into the Perl code. Perl is a programming language. It's very old, very useful. It's not really strongly… Like Pascal forces you to do things – you can do whatever you want in different ways. And if I define two variables; one numeric, one character-based and then I define an array with colors.

In order to make the arrays in color you need to basically have each registrar have its own color. Now, I didn't want it to become really complicated and start mathematics to calculate a color table, I just used a number of colors here. Here you see eight with the ellipses in the middle so it's more than that. And if I find one that doesn't display well I just remove it off the array.

Then Perl has the structure called hash, which is similar to an array. An array you can access in the first or second or third but a hash is similar to that but you can access it by keying in. And then in the Perl program we set up an object called Graph where we define the pen width with the thickness of the arrows that we're using. One could have done even more detail.

For example at the program [Run? 00:24:31] look at if it's a big volume we make bigger pen width and things like this, but this complicates it too much so I just did some simple stuff to look at it. Then we connect to the database. I left the password out because you don't really need to know that. And we used Tunnel – as the system sits on a separate computer we access it through Tunnel, the local host, so that it is secure.

We construct an SQL query. Here we look at transfer requests as they're approved and we use only those registrars that had more than five transfer requests. We assign it to a variable, which we then, at your own peril, execute. Since you go into a production database you have to be really confident that what you're doing is correct because you can make a mistake.

CoCCATools, which we use, has the advantage that to make any simply change you have to first set some trigger conditions and otherwise if you do just a select statement and pull the data off you cannot change the existing data. Basically you get a table like this. Interesting here is dic lost nine to ihostnamibia and ten to Verizon. Verizon on the bottom lost six to intertech and 33 to itn. Of course we've got about 30 registrars; I just showed you two.

There are some registrars who only gained and some registrars who only lost and some registrars who gained and lost. And then you basically access this table when it comes back. But you must first of all draw all the edges, which means the lose from one registrar to another. But I don't want to change the color for each registrar.

I want each registrar to remain the same color so basically we run this, we look… Did I look at this before? It's sorted. This is the second line of the same registrar and then we don't change the color. Otherwise we change the color. And since I only have a certain number of colors I look at the first line, which says it is bigger than $#colors. It counts the number of colors in that area that I showed in the second slide.

If I remove one I don't have to remember to change the numbers. Once it reaches the end it reverts to the first so you recycle. If you have 30 colors you can take even more than 30 registrars but at 31 we'll have the first color again. Then we put these edges into the object and then we go through this hash and only look for each registrar once. Okay?

And then we basically output this to an .svg file – scalable vector graphics; this is a vector image – it's standout, it's an .xml file format.

And you can also output it if you want into the language that I used to do this, which is Graphviz. This is the output in the language. It's very complicated but you don't really need to manipulate in manually. When you run the program on it it produces this.

Here you see all registrars that have lost more than five domain names. The interesting thing here is that the escrow domain, which is the third on the left and the second row from the top lost a few to gijima, lost a few to swakopcom. Swakopcom lost a few to pyxix, who lost a few to dic. Pyxix lost a few to gijima. Itn in the middle lost a few to pyxix. And these other ones from this one reseller, who moved from one registrar to the other.

So you can actually follow this very nicely. Now, if I do the same for more than one transfer you'll see the international registrars all classed on the bottom left. They basically don't interfere with the Namibian registrars. We have a policy that foreign clients can go to Namibian registrars but foreign clients pay more than Namibian clients.

So they usually tend to go to foreign registrants and when you do that… The software automatically – with just a little bit of cheating from my side – classed this all very nicely. This is just the foreign registrars, the ones that were on the bottom left. And you can see, even if you enhance this picture, the graphics are very smooth.

So if you want to do other things that is some software that I think we can use. It's a different process but we use the same data. The top half is the Namibian registrar, the bottom ones are the foreign ones. It's not really helpful but there are, as far as networking is concerned, circular structures where this simply process may become helpful.

Then we had another little problem. In 2005 the request by local resellers to create a registration system that registers [inaudible 00:29:55] was a big issue for them that we were not able to take registration 24/7. Of course, the registration science for those resellers do not allow registration 24/7 but they said you must do this. So as you know we then deployed CoCCATools who staffed this. The resellers became registrars and therefore we can look at any data. And for us of course the question is, do they now that they can actually do it do it 24/7?

R is a programming language which is particularly strong and designed to do statistical analysis. It has got lots of models, graphics, it can interface with LaTeX, you can write reports – it's very good. Again, we use the SSH Tunnel. So I'm showing this because values are not assigned with an equal, it's a diverse arrow. I don't really know why this is but it is the way it is.

You can also use an [eco sign? 00:30:53], a font [inaudible 00:30:54] recently, but I didn't know that until now so for many years I always used the data. So you look for one year, that's 2012 until the end of May 2013, because that's the report we generated for this Commission. And also it's got a postscript module and it has got a plotting module.

So we connected with the driver and then we create a query. And here we look at registrations and renewal so the trans type in the middle – it says "trans_type ilike 'Re%'" catch captures registrations and renewal.

And it's date constrained and we only capture the hour of the day of the timestamp that we have on our database.

And then we group it so we get for 24 hours, for each hour we get the number of registrations per hour. Then we do some housekeeping processing. The red query is only the one for the Namibians. We also have to do it separately for the foreign registrants. We have to count the number of registrations, we have to… For the drawing of the axis, if you don't do the last line of this slide you will get a number that is exactly 375, which is the maximum.

It doesn't look nice so I tried to round it to the nearest 50. Then you plot the lines, you plot the lines, you plot the bars with a different command and then you do the things around and that is how it looks. It shows that the Namibian registrars are unionized and they start working at 8 o'clock. And they go for lunch at 12:00 and start at 2:00 and then they slack off at 5 o'clock and that's it.

We can also see owner operators working late until 10 o'clock in the evening. The green lines are the foreign ones. We have a few from Asia and from Australia so you can see them on the left because of the time zone. Then we have a significant number of European registrars – you can see them in the middle and the early bird Americans also. And then they slack off in the evening when the time zone goes to the Pacific side of the US.

The number of transactions is listed in the box on the top right. It's just for one year. It shows clearly that a request made by this registrar was not really relevant and was not really honest as far as I'm concerned.

We wanted to see how we could display… We have domains in the top level and the second levels so we wanted to display the percentages.

So basically you figure out how to construct a query. Here I counted the numbers grouped by zone but I also counted percentages. And of course the name of the zone. And then I used this 3D plot page which then looks like this. If you look carefully the legend is not really high resolution because at that stage I hadn't figured out .svg for this particular software. I will try and play with this to get a high resolution of it.

However, what it shows is that 72% of our domain name registrations for that year were in the com.na domain, which is very well branded. Our local Namibian registrars and clients in particular don't really want top levels, they want .com.na – that's what everybody wants.

Also, a large complement of .com and for historic reasons the [inaudible 00:34:55] part of our business is German-speaking so they have got German clients so they use the [.eu? 00:35:08] because they think it's easier. But this is the thing.

On a little side note, since .co.na is approximately 12% of all our domain registrations in that year, roughly 250, you would probably be surprised if I told you that somebody has put up .co.na for auction at [Zidu? 00:35:35] and it wasn't me.

As a totally separate issue I exchanged some nice words with the German-speaking representative of them here, but we are busy taking this up with them because there is something seriously wrong in their system when a generic top-level domain is being put up for auction by somebody who doesn't own it.

It's not a name like lisse.na or medina.co.na, it's got 200 domain names and somebody is trying to sell it. Something is wrong. Okay. Of course, much more is possible. This is simple stuff, I just wanted to draw some nice pictures to make some points during this presentation, but the more you do this and the more [hiding? 00:36:23] you get when you dare to ask a question on the Usenet for the German-speaking Usenet forum. The more you start looking into this you find out that there is cool stuff available out there.

Things that can be done at SQL can be done in FRED as well as CoCCATOOLS. .na uses CoCCATools. Some African countries and some overseas countries use the excellent FRED package that was developed by [inaudible 00:36:50]. Also uses post-SQL so the tables are a little bit different but you will probably figure out pretty quickly what to do if you wanted to take this.

As I said, it is a concern when you go into a production database but in CoCCATools, if you do only SELECT, if you don't try to make any modifications you cannot actually make any changes. LaTeX is typesetting software that is not like in Word, where you see what you type, but it's more like a programming language.

But the point here is that you can use this tool called SWEAVE in R, where you put chunks of the code into a LaTeX document or chunks of LaTeX code into the R document and then when you run it for a certain period of time it does the queries and then constructs the tables into a graphically appealing form. It does the graphics and puts them out, so you don't have to rewrite your report when your data gets updated; it updates it automatically.

Same here for this presentation. This presentation is also written in LaTeX. I used a font and [inaudible 00:38:05]. I used the Beamer tools. I did not use SWEAVE because I hadn't figured that out at that stage, but still, if I didn't put a date constraint in, for example 31st of May; if I was to look it up today I could have just run it again and it would have automatically updated the images. And that's kind of cool.

I don't think with a Microsoft PowerPoint or Open Office or Keynote can do that. I think when you paste the pictures in there they are in there and you have to manually take them out and put an updated one in. And while I'm being so epic and going into so much detail about it, I specifically wanted to show that with open-source tools you can produce very high quality, good-looking output very easily and without much work.

Any questions? Come on now, I didn't want to dazzle you I just wanted to make a little presentation. Okay, has Tom Barrett in the meantime arrived? There you are. Excellent. [pause] I apologise for the little technical hiccups but we had to ask the PDFs to be sent in the day before so we could upload it for the remote participants. Kristina, please raise your hand if there are remote questions because I tend to get ahead of myself here.

TOM BARRETT: Good morning everybody. My name is Tom Barrett and I'm going to talk today about how country code registries can leverage the Trademark Clearinghouse. As a way of background I'm the tech contact for the .pw registry. I also have an ICANN registrar. But we have a third business called tm.biz that does trademark validation and so we recently completed the .PW Sunrise and we also did the radio fm and radio am sunrise periods.

And these sunrise periods were done in the very traditional manner where corporate registrars asked their trademark clients for trademark data. That data was provided in files and then we had to verify that data against online trademark databases. That type of model probably is soon going to be obsolete and if any country code is thinking about running a sunrise period you should probably consider using the TMCH instead.

As more and more trademark owners submit their data to the Clearinghouse it actually becomes a very efficient mechanism for any country code registry that wants to undertake a sunrise period. Just as a way of background, what is the TMCH?

Obviously many of you may know this, but it was created as a repository of validated trademark rights, and it's either registered trademarks of a national jurisdiction, trademarks awarded by court order or trademarks awarded by international treaty. And ICANN has already hired vendors to run it for the initial round of TLDs. Deloitte is called the validator and IBM runs the back-end database.

And it's noteworthy to point out that Deloitte will eventually have competition and there will be multiple validators but IBM will always have the sole responsibility of the back-end. So the process is very simple today. Trademark owners can sign up for an account directly with the Clearinghouse or they can go through a trademark agent.

They submit their trademark data and Deloitte will then validate that data. They may also provide what's called Proof of Use and if they take that optional step then they will also be issued with a token called the SMD file. A quick note on trademark claims. It is that 90-day period during the initial launch of a TLD. Every trademark that is accepted by Deloitte for the Clearinghouse is automatically opted in for the claims process.

There's a two-part feature to the claims process. Any registrant going to a registrar requesting to register a domain name will be shown a note saying: "There is a claim against your requested domain name; there is one or more trademark owners who have this exact string as a trademark, are you sure you want to proceed based on this information?"

If they say yes and proceed that name not only gets registered but then the registry will inform the Clearinghouse about the event and the actual trademark owner matching that domain name receives an email. The second use of the Clearinghouse is for the sunrise period and this is where trademark owners provide a Proof of Use component.

And they are provided with this SMD file, which some simply call a sunrise token, and they have to provide that to the registrar at that point of registration during the sunrise period. So that's how the Clearinghouse works today. Let's talk a little bit about what this might mean for ccTLDs. This sunrise token that I talked about, SMD stands for signed marked data.

It's basically a file that is encrypted. It has a digital signature encrypting the file. It basically contains all the trademark information that was submitted to the Clearinghouse, the trademark owner information, things like the registration number, the jurisdiction, the goods and services, the classes. If there is an agent involved in submitting to the Clearinghouse it would contain that.

It also includes all of the domain names that the Clearinghouse determined were an exact match of this particular trademark and there could be multiple domain names because a trademark could have special characters that have to be replaced or omitted. They also are supporting this idea of variants that were the result of a successful UDRP.

In terms of the SMD file, what makes this unique and reusable for other registries other than New gTLDs is that it's literally a file – it's portable. It stands on its own so it can be used by any registry, such as a ccTLD that wants to run a sunrise period.

In the future the type of sunrise that we just went through for .pw, where you go to trademark owners and ask you to give you trademark data – that model is probably going away. Unless you have a very specific trademark validation requirement such as geographic territory or an industry that the Clearinghouse would not be able to support.

So instead the trademark owner is simply going to, for every New gTLD, provide this sunrise token in order to do their sunrise registrations. There are two ways that ccTLDs can make use of the TMCH. First of all for sunrise there is this SMD file – it's portable. All they need is the public key that will be managed by ICANN and IBM. They can read the digital signature of that file, they can determine that the SMD file is not on a revocation list and then they can basically parse the SML within the file and validate any sunrise registrations that might exist.

So this is good news, very simply, for ccTLDs who are contemplating some sort of trademark sunrise period. Now, on the claims side, this is a new service coming from the TMCH, specifically for ccTLDs. They're going to provide a trademark claims service as well but they're going to roll it out in three different, distinct phases.

The first phase is simply going to be notices to the trademark owner. If you're an existing ccTLD, you've been up and running for 10 or 15 years, you could implement this additional rights' protection mechanism as a way of branding your TLD as trademark owner friendly.

And simply every time there is a registration there is a way to query the Clearinghouse, see if that registration is an exact match against an existing trademark in the Clearinghouse and then the Clearinghouse would notify the trademark owner that that name has been registered.

That's phase one of this service and I think this availability is probably later this fall. Phase two will be the other component of claims, where a registrar for the TLD, when a registrar attempts to register a domain name, will be able to determine if a claim exists for that proposed domain name. And it would simply be a yes or no.

It won't contain any trademark information but yet it will be more information than a registrant might have today and the registrar can communicate this fact to the registrant and ask them if they want to stop and do some research or if they want to go ahead and proceed with their registration.

And phase three will be a full-blown trademark claims service from the Clearinghouse, where the potential registrant is not only displaying that a claim exists but they will also be told more information about the background of the trademark owners who have those claims. They'll be asked whether or not they want to proceed and if they do then eventually the trademark owner would get notified of that.

So this is fairly new information but it's an interesting way for ccTLDs to leverage the Clearinghouse going forward. Any questions?

EBERHARD LISSE: Thank you very much. As an aside, if I'm not mistaken the CoCCATools software has got interface to the TMCH directly, so should you use that and you wish to use it it's already connected. Questions from the floor? We've got two microphones here but we can also go to the seats. Are there any questions from the remote? Thank you very much Tom.

I can't see Pierre from here. Pierre Dandjinou, are you here? Fascinating. Okay. So next would be Ben Fuller. Ben is a Director of Namibian Advocate Information Center, which is the .na ccTLD Manager.

He also is the Dean of the Faculty of Sustainable Development and therefore has a different perspective on technology and things than I have.

Given that we have an African context here, and especially with the African Initiative looking at increasing the number of registrants we wanted to share a little bit of what we think about these things.

BEN FULLER: Good morning. As Eberhard says, I'm a bit sad that Pierre is not here yet because it would have been nice to hear what he had to say before I gave my presentation. But we're talking about how we promote DNS in a developing country. Why are we interested in development?

First of all, both Dr. Lisse and I predate Namibia's independence. We were both living and working in Namibia back in the 1980s. He was a medical officer in Oshakati, which is in north-central Namibia along the Angolan border. I was an anthropologist running around on the edge of the Namib desert about 600km to the west of him.

In 1995 he and I connected in Namibia through the Internet. Now, our motivation at that time was of course post-apartheid development, because that was very important in Namibia. As I say, both Eberhard and I had seen Namibia under apartheid times and we were very concerned about how we could develop the country.

We now have the ICANN Initiative for Africa and we feel that we've got a few things after 25 years of developing the Internet. We've got some insights and ideas, we have a sense of what will and will not work. Now, for those of you who don't know, Namibia is a country of twos. It's twice the size of California. It has 2.2 million people. We have a lot of space. Never trust a Namibian when they say, "Oh, the drive is not far."

Far can mean 100km. So anything below that can be "not far". So if you're driving around with a Namibian and they say, "It's just around the corner, we'll get there," you might find yourself driving long into the night. It's also two countries; a rich and a poor country. This comes out of our apartheid and colonial past.

The minority of people lives a lifestyle similar to that of Canada or Sweden. The vast majority live a life closer to that of people living in the Sudan or the Congo, but it is getting better. Namibian development over the last 20 years has seen some moderate successes. Poverty has been decreased by 45%. Primary school attendance is up to 95%. Namibia has become an upper-middle income country. We're no longer technically a developing country and we have a growing middle class.

Now, as part of this we have also gone through the mobile revolution that has taken place in Namibia since about the late 1990s onward. I was saying in one of my classes a few months ago, I was telling my students about the bad old days when you had to have a phone that was connected to a wire that connected to the wall and about half the class looked at me and said: "No, that didn't happen, did it?"

So we're going through the rest of Africa the mobile revolution. We have roughly 96 mobile phone subscriptions per 100 people. Now, Namibia includes people such as the San bushmen that are very famous, the Himbas, the [chimbo? 00:56:45], the [Oluthamba? 00:56:47], who are almost like Africa's Amish – they live an old, traditional lifestyle, but if you go to their population centers around where they stay you'll see these guys pulling out their cell phones and checking things out.

Roughly 90% of the country has mobile network coverage, which is actually quite considerable because it's a very large space with a very small number of people. We have 260,000 Internet users, however probably 60% of them are using the Internet because of Facebook on their mobile phones. They may not even know that they're using the Internet, they're just more interested in Facebook and what's happening with their friends, boyfriends and girlfriends and that.

So we have to talk about who we want to reach with our DNS development. Do we want to deal with our Canadians and our Swedes, or our emerging middle class or the Congolese and the Sudanese? And the ideal is to reach all three. But you have to understand you need different modalities for reaching them, however we've realized a long time ago the key is registrars.

We run the register for .na. Registrars deal with the clients, they do the marketing, they deal with the general public and that's very important. We have a number of ICANN-accredited registrars. Some things about them that are important is they use debit and credit cards – you log onto their websites, you put in your credit card number, you get your domain name, everything's fine, their language is in English.

They have a very limited local presence. In other words, many Namibians like to actually go in and talk to someone; they're not interested in… They're a little bit leery at the moment about logging onto something and giving some cash to somebody you don't physically know.

Then we have a Namibian-accredited registrars that are able to, for example, take cash payments. They will speak in the local language. Now, this is very important because many of our emerging economic middle class and the emerging SME owners and other people – for them the Internet is new. Even though they may speak English, they want to hear about it in their local language because that's what they're comfortable with.

And by having Namibian-accredited registrars they can understand these concepts first in their local language and then understand it in English. And we also have… It's good to have a community presence. As I say, when most people are approaching a new technology; something very new, such as the Internet, they like to come in and actually talk to a person. Not a chat room or anything like that.

So for us, when we're setting up, how do we reach these three different markets that we have? We have to think about the kind of registrant and what their capacities are. It's vey easy for MarkMonitor and Verisign and others to connect to us; they just log in, get an EPP connection and everything is fine.

For some of the younger people we have web developers, we have small ISPs, just people who are trying to provide some sort of Internet service and might have 10 or 15 clients and domain names that they register. So the thing we like about CoCCATools is that it provides a wide array of options for financial and operation aspects. It gives us good data security.

It's got a fairly easy learning curve. We train all of our local registrars – it's about a 45-minute course in how to use CoCCATools. CoCCATools is very intuitive and it's got scalability because a registrar has room to grow. If we get a good active registrar, we can for example give all of them credit facilities so that they don't have to run out of money if they're registering a large number of domain names.

We have a facility whereby we can allow… Registrars are able to begin to start learning EPP and logging on and ultimately transforming from a web interface over to EPP. We just had one web registrar do that. So we ended up choosing CoCCATools as our platform, because you need a lot of options in a developing economy, because you have this multi-tiered level of interest and skills in terms of connecting to the Internet.

Some useful policies that we've found over the years that help develop the Internet market. First is that foreign funds locals. Setting up a good, high-class register is expensive and you need a fair amount of money. So we came up with a two-tiered pricing structure. Non-Namibians pay more, Namibians pay less.

If you come to Namibia and come to any of our major lodges or major parks such as the Etosha Park, or if some of you leave here and drive up to Kruger National Park in the next few days you'll find with your American or Swiss or German passport, when you show up, you're going to pay a lot more than someone with a South African passport. This is a very common pricing structure here in Africa.

The next is registrar accreditation. We provide minimal checks on registrar accreditation. However, in terms of recent legislation about the use of money laundering and other sorts of illegal activities using accounts, we've started to be a bit more rigorous about knowing our client, trying to get peoples' founding statements, getting some background documents on people.

And lastly, we make everybody sign a contract and the contract specifies very clearly what registrars can and cannot do. And it's very important that you hold them to your contract. We insist on WHOIS accuracy. For us, partly due to legal issues that are taking place in Namibia, the admin contact is a very important part of WHOIS accuracy.

For us it has to be a real person with a real email address. This promotes the integrity of the ccTLD because if a domain name is being used for phishing or other activities, somebody can find that person involved.

Lastly, something that we feel is very important for a developing market is transfers between registrars. Registrars like to make themselves the admin contact of someone's domain name and then not allow someone to transfer from one to the other if a registrant feels they have suddenly found a better deal or someone was provided a better service.

So we feel that transfers between registrars are free. Nobody is allowed to charge anything. We don't allow to charge anything and it is written into our contract that a registrar cannot prevent a transfer if there are any commercial considerations or payment disputes between the registrant and a specific registrar. We tell them to sort it out amongst themselves or we will do that for them and they will not be happy.

We feel that this is how you promote competition. Running a registry, we can't promote competition in all the market, but this is how our little part of the process can promote competition. And I think that's very important for developing the local, smaller registrars who deal with the middle class and the Congolese and Sudanese aspects of our population.

You've seen this picture already but just to show you – this is what you get in our local market when you start transferring registrants around. And it's very useful because it keeps the local registrars on their toes. They have to take care of their clients because somebody can move somewhere else very quickly. Okay?

Lastly, a few things about DNS and development. I've been working in development in Africa for 30 years so there might be a few things we want to take forward when we're talking about the ICANN Initiative in Africa. First of all, DNS is not a universal panacea. DNS development will not solve the digital divide but there will be people who will think that it will.

We can play a part in a larger process of ICT development, but ICT development includes dealing with things like infrastructure, laying fiber cable, connecting to the under-sea cables that are running around Africa now. Legislation – do you have a Telecoms monopoly in a given country? What is the market structure like? As I've shown you, not all African markets are the same. What are the national and local policies on how to develop DNS and run a ccTLD?

And lastly, you have to understand African markets on a case-by-case basis. Africa has over 50 countries and it ranges from Islamic-dominated countries up in North Africa and down here to South Africa and a lot of stuff in-between. There are many different languages, legal structures, colonial legacies, banking systems, foreign exchange regulations. Every African market has to be understood on its own terms. There is no universal cookie-cutter that gives you an idea of what's happening here in Africa. Any questions?

EBERHARD LISSE: Come on, there must be some questions. Are there any Africans around here? At least three sitting on that table…? Anyway, Pierre has arrived. You can sit on this chair. We've got your laptop organized. [background chatter] We're having technical problems. Just hang on for a few seconds. [pause]

PIERRE DANDJINOU: Thank you very much and good morning to everyone. I'm sorry I was late, I was caught up in other meetings. I'm happy to be here and to share what we've been trying to do as far as the Africa strategy is concerned. I'm supposed to take you through a few slides but I will just start and then once we are ready you'll have the slides.

I think the Africa strategy came as a tool of further engagement of ICANN and the African continent. I think one of the things that really triggered the development of this strategy was the combination that [inaudible 01:11:50] with the New gTLD program that there were not many applications coming from Africa. The very few that came from

Africa were actually coming from one country basically, so there's something that needs to be done.

But also we are aware of the way our different ccTLDs have been functioning in Africa. There are issues with them an in fact the [inaudible 01:12:29] Forum we had here on the DNS really showed some of the issues again. So the Africa strategy was supposed to be a tool for engagement with the African continent.

The idea was to really see exactly what's to be done to really change the landscape in Africa. There was a Steering Committee that was set up and the Steering Committee did its surveys and discussed with the community. And a few things came after that that were definitely [inaudible 01:13:10] as a market.

And I really like what the previous presenter was saying; that you need to understand all the African markets and he was also saying that it has to be on a case-by-case basis. We don't have one Africa; we have 54 countries and so the realities are not the same. So it's important then to actually see Africa as a market.

So this strategy definitely welcomed that and came up with strategic objectives. This [game is a? 01:13:52] new season because that was the first time that ICANN wanted to really engage this continent and also… I was supposed to go through the policy-making process but I want to skip that and go straight to the African strategy.

This strategy actually has eight objectives, which have been further developed into projects and that really also have derived some indices on measuring progress. And of course an action plan has been attached to that and therefore it becomes monitoring, evaluation tools also. It's a three-year plan. We've actually started since January to implement some of the [inaudible 01:15:00] priority project.

One of the priority projects that we felt was to make sure that African ccs and registrars offer a secure environment. One of the things that was done was a DNSSEC roadshow. We initially selected eight countries and it has started and we are also trying to grow this number and starting for the financial year 14 to have 16 of these countries instead of the current eight that we have.

Now, the other project is about the… I will skip this. The other project is about the whole idea of outreaching to the registrars and the registries; how do we do that? So one of the ideas was to organize a governing of all of those and we did this in Addis Ababa and we have this workshop, which is quite interesting because I actually had an impression that the registrars and the registries were actually talking in Africa.

The good thing there again was to see some of the issues that were highlighted from this meeting. And then this also helped us see exactly what we should be doing. A project we'd like to also do within this framework is also the whole thing about the observatory domain name in Africa; afTLD with ISOC, for instance, now are ready to do this and we will be partnering to see how we move on that as well.

Obviously there is going to be a study of the African market of domain names and we are actually calling for other partners, because actually we really need to know what the realities are on the ground. So that's what these studies are going to do. In the next six months we'd really like to launch this study as well. But the DNS market in Africa, if you… For instance, one of the CEOs of ICANN was in [that as well? 01:17:23] we need to multiply the number of registrars that we do have from five, and we'd like to be moving to 20.

But of course, for that we also need to make sure that the market is there. The registrars we have today are not really making any money per se, and so one of the issues they are raising is for instance the level of insurance they have to pay and so that's taking all of their money.

We are working to see how [inaudible 01:17:54] within ICANN and also with insurance companies, how one could really lower this barrier. And that's something we want to push forward. One of the things I'm also noticing that we will try to do is make sure that the global registrar leaders also understand the issue in Africa and also helps develop Africa as a market.

That's why we are happy to have most of them at this DNS Forum. And then we did a change and we are going to have some specific [inaudible 01:18:34] with them. As far as the African strategy is concerned we are ready to… There was this idea floating around that we need incubators for young entrepreneurs.

There was this idea to have exchange programs where you might have a few Africans reside for a few weeks or months with DNS leaders. That's fine, that's something we'd like to proceed with. But I think the most important part of it is also that we do have these collaborations on different regional levels. There are issues pertaining to policy development but there are also issues pertaining to the legal environment that can facilitate this market's development

We are looking forward to ideas on those factors and to make sure that the DNS business thrives in Africa. Finally, I think the action plan incorporates some other issues in terms of outreaching to, for instance, governments in Africa, because this is also an issue people are facing. Cctlds now…

Well, we're having ten countries at least requesting a redelegation. They are putting questions on how the ccTLD is being run at home. Some of them haven't decreed; they say, "Okay, it's for us to manage," so there is some debate maybe that they need to be conducted on country levels totally to be [consensus? 01:20:17].

So all of these are real issues that we need to tackle. So basically, that's what the Africa strategy is about. It's building capacity within African ccTLDs and registrars and seeing how we can really make Africa a real market for business. So in a nutshell that's what it's about. Thank you for your attention and I'm ready for any questions if you have any.

EBERHARD LISSE: Any questions? No. Okay. Next one would be Nishal Goburhan. I don't know him but he is here? Okay, good. No, we haven't met yet but that's why we communicate via email. Nishal works at AfriNIC, the regional address registry and they have been doing work on fostering Internet exchanges.

For example, in Namibia there are really big problems because the ISPs don't trust each other. They don't want to do this for whatever reasons but it is apparently through the support of AfriNIC coming into many countries. They are also are involved in deployment of the L-root that is under ICANN's control. ICANN has an initiative and the regional registries are involved there.

And then they have been pushing DNSSEC with the DNS roadshow, so we asked them to give us a bit of an insight there.

NISHAL GOBURDHAN: Thank you Dr. Lisse. Good afternoon ladies and gentlemen. The topic that you see on the Agenda for today is a slight misnomer. Apart from spelling my name wrong it also speaks a little bit about L-root. We are at an ICANN meeting so it would not really be appropriate for me to talk about L-root. Instead, as the kind doctor has indicated, I'll talk about the AfriNIC plans that we have and how we get L and the other root operators that we work with connected into our regions.

This is my first ICANN meeting. I have a green badge. It's good to be here. My name is Nishal and as Dr. Lisse said, I work at the Regional Internet Registry for Africa. My job is pretty much to be able to work with our region, the Africa region, to try to help operators in the region. That could literally be a Telco operator, a DNS operator or anybody that's in the area understand and improve the infrastructure.

So my bias against a lot of what you're going to hear today is how do I build infrastructure to get things done? Some background: I spent about

12 years working at an Internet service provider in South Africa. From there I moved across to AfriNIC and I do infrastructure work there, and I also run the exchange points in South Africa. So if I say something and it sounds like I am taking sides, please understand why.

We'll start off with just talking a little bit about what AfriNIC is doing in terms of helping Internet exchange points in Africa become sustainable and grow, really. A little later on in my slides I have a map of where exchange points in Africa are. And one of the problems you'll hear, not just from me but from anybody, including good people like [inaudible 01:24:44] in the crowd is that there is obviously not enough Internet exchange points in Africa.

So what AfriNIC is doing… Specifically, the first thing I always make a point of saying is we treat Internet exchange points as critical Internet infrastructure. Now, if you're at an ICANN meeting you already know about bottom-up policy development process so I'm not going to explain how all of that works through us.

But through that policy process that we have, Internet number resources, essentially the element, the thing that AfriNIC has to offer any operator in the space is made freely available to all Internet exchange points in Africa. You'd think this is an easy thing, this is a small thing, this is not really something that exchange point operators should worry about.

But the big difficulty that those of you running exchange points or who've worked with them will know, is that it's really difficult when an exchange point is running to get peers at the exchange point numbered or to get things done properly. So we're trying to make it upfront and clear to them from the start we have these resources, we'll show you how to use them, we want you to use them and yes, they're free.

So if you're an exchange point operator sitting in the room, in Africa, or you're starting up an Internet exchange point – and I can see one of you here – please speak with me a little later and I'll help you get the number resources you need to make that Internet exchange point run properly.

I will start by saying we don't do box-drops. I used to work at a company a long time ago that used to just drop black boxes on site and the black boxes will fix things for you and then the company makes a lot of money by coming back and fixing these black boxes when they don't work. Well, we have a very different approach to building stuff and our goal here, our belief system, is sustainability.

So for us, an exchange point is more than just a simple layer two switch. There are about 14 Internet exchange points in Africa that operate on that basis. It is a simple layer two switch that's being plugged somewhere into some closet – you would find it hard to believe – where somebody was shown how to interconnect maybe ten years ago and that's it. Everyone's scared to touch it. Everyone's scared to move it.

And our philosophy behind that is in order for an exchange point to succeed, growth and sustainability are key success elements for the exchange point's success. So don't ask us for hardware. I know that sounds strange, but I don't work for a switch vendor. There are some other good people in this [car? 01:27:29] that have rented equipment to an exchange point. We can't do that.

What we do is we help you donate and we work with people like the Internet Society ISOC or the Access Project. We help you develop policy, we help you develop guidelines, we help you develop technical ability to run and build your own Internet exchange point. I already manage three; I don't have time to do any more and nobody else in our environment does either.

So we really want you to become self-sufficient and to get your involved in managing your own resources, in managing your own environment. And that's what AfriNIC and our partner in this, ISOC, are aiming to do. So far, with the Access Project – which is an AUC-funded project that's been running since about October last year – AfriNIC has supported ISOC and facilitated five technical workshops and ten TA policy workshops.

And that's since about October 12th last year. And this is an ongoing project so we hope it will continue. You would have heard the Doctor say this morning that there was about 33 countries they're aiming to get into so hopefully we'll have all of this covered. Attached on this already, what we do offer is free – free! – expert assistance with sustainability rebuilds.

We're doing this right now in one SADC country: we're helping an existing exchange point that was put in by some very expensive world bank consultant many years ago, to re-architect and design something that will actually succeed and work for the future. So we're moving them from what the techies would call the closed layer three exchange point to a simple model that we know works for us and we've seen working in some of the largest Internet exchange points in the world. So that's one of the current projects that we are busy with.

So that's just a little bit about the work we do, particularly around Internet exchange points. And again, the goal here is simply to enable our constituencies, enable the environment that we're working in. And this is the African scope. My work in AfriNIC more formally relates to the DNS projects that we run.

AfriNIC runs three separate DNS projects and as part of our goal, as part of what we are expected to do to support community in Africa, we have developed these three projects that are available to the community – again at no cost. So if you're in the African region, if you have operations in the African region; if you're a ccTLD Administrator or something then some of this is probably particularly interesting to you.

The first project is called the Africa root server copy project and for short – because we like to do this – we call it the AfRSCP. And the goal of the AfRSCP is really to try to get local copies of the DNS root servers into every country. So that was the ambitious goal. We started the work for this project in about 2009 and the first DNS root server was a K-root that went into Tanzania and since then we've been slowly deploying stuff into Africa.

Our goal really is to be able to say that we want to make it easier; we want to improve the level of Internet infrastructure in our region. Obviously you've heard all of this before and there are other people who are doing good work in this field; ICANN on its own with its L-root project is doing stuff like this. [Autonomic and ? 01:31:14] F-root have multiple copies of [DNS nodes? 01:31:15] in Africa.

We were specifically targeting at the time areas that have Internet exchange points. So when we launched the project back in 2008, 2009, literally we thought: “Great, we can help develop the Internet exchange point and help develop the DNS infrastructure in the country at pretty much the same time. So spot the problem.

The green triangles that you see in the map are the countries that have know working Internet exchange points. That means they’ve passed some traffic in the last couple of years. Now, some exchange points are a lot more active than others. I can speak for the South African ones – they’re vibrant, they’re active, they actually do a lot of traffic.

But some of the smaller Internet exchange points, sometimes, even though they have participants at the exchange points they suffer from problems where the local operators don’t necessarily advertise all their networks at the exchange points locally, for instance. So they’re connected but they’re not actually seeing the benefit of it.

And the problem that we ran into practically as AfriNIC was that when we tried to get to these exchange point operators and say: “Knock knock! Hello? We have this project we’d like to work with you on. We want to get copies of your DNS root server into your region,” was really difficult.

Because again, most exchange points were being treated as a simple layer two switch stuck in a cabinet somewhere and there was no… I want to use the word business but I want to use that word loosely. There was no business built around it. There was no sustainability built around that. So the project that we had specifically was meant to deal with invigorating exchange points. And the way that… We were actually given a nice way around that thanks to ICANN.

In October last year – that's my CEO, [Adieou? 01:33:07] – ICANN an AfriNIC signed an agreement where AfriNIC would work with ICANN to help improve the spread of particular L-root servers across Africa. So initially we had three DNS roots on board. We had I, we had F, we had K – and K has recently changed its policy – and now we also have L on board.

And since then we've been installing equipment, we've been working with ICANN and obviously the other roots to get things in place and to get equipment installed. So we've got installations that have gone in [inaudible 01:33:44]. We have hardware on the ground. A really good friend of mine says he is an engineer and he judges by results.

So we have got equipment on the ground in three countries right now. We are just waiting for a few logistic issues at the exchange point to be sorted out. We are at advanced stages of completion. In other words we have signed agreements. People that are saying: "We want to work with you, this is where the equipment is going to be housed, these are responsible people that will be working with you on the project."

So we're making slow but steady progress. And for me what's important is I can see that it's working in getting these exchange points and these local environments sustainable. So it's not just – as I said earlier – it's not just us dropping boxes. If you're an African operator, if you're from an African country, if there's no DNS root in your country or you want to help improve the state of DNS there, rootproject@afrinic.net comes through to a ticket system that I read.

We are happily accepting new applicants. There is always paperwork, unfortunately, but we'd love to work with you to help improve the state of the network in your environment. So that was the first of three projects.

The second one is the RFC 5855 supported servers, and that's fancy talk for the servers that are responsible for IN-ADDR.ARPA and IP6 ARPA. And every since 2010 these would be delegated to the RIRs and ICANN. So the six of us effectively provide name services for this. We've been running IP6 ARPA since 2010 and IN-ADDR.ARPA or the C-version of that since about 2011.

And of course no technical presentation is every really a real technical presentation if there are no statistics involved. So here is a snapshot of some of the statistics we see, just for IPv6. Bear in mind that this is caught in Johannesburg so it's at the bottom of the world, which is probably where most peoples' resolvers are not getting to.

We do about 1,500 queries a second and for IN-ADDR.ARPA we do about 3,500 queries a second, so that's just… So why is this relevant? I guess the reason I'm telling you this or telling the Africans in particular here is because AfriNIC is now running the CDOT server and what we are able to do, because we have control of this, is through localized anycast programs we're able to get this into your environment.

So again, working towards making – at least the DNS, which is something we all care a lot about – the DNS in your environment a lot more resilient and performing a lot better for you. So it's the second project that we run that we can bring to you and help improve the state of connectivity in your region indirectly but directly as well.

And the last project that we run is the African DNS support program or the AFDSP. AfriNIC was really small. Just a little bit about it: in 2009 when I joined there was probably one router and main network and just three servers at that stage. And since then we've been growing our infrastructure to support a whole bunch of other projects that we're doing.

Obviously you know in today's modern age the side effect of building infrastructure is that you always have stuff lying around; spare capacity as I like to call it. So through the effort of AfriNIC building its own cloud environment we've got the spare capacity available in our infrastructure. And one of the things that we've been doing with that is to make the spare capacity available to holders of critical Internet resources.

So if you're an exchange point and you have resources and you want to reverse DNS, or African ccTLDs – we'll make this available to you to bolster your infrastructure. Now, it's free servers. We don't charge anything for it. Again, it's part of the service we provide to our community. We want to make you infrastructure resilient and that's the goal, that's the guideline of this entire project.

And we can make this available either across our infrastructure or part of it; whatever it is you choose. Now, there are already people doing this. There are lots of really good DNS operators that manage this. But as you heard earlier today the focus is on African solutions so this is an African solution. We manage this stuff. It doesn't leave our area of control. If you've been reading the news you shouldn't have those worries.

We currently support – I double-checked this morning – 17 African ccTLDs. We're in the advanced stages of set-up, which means we've been speaking to the people, they've responded to us, I've got the information, we've been verifying PGP keys or in some sort of technical set-up. So we're working with another four operators as well.

If you're interested, if you're an African ccTLD operator there is a sign-up page on our website or I'm here the whole week along with my colleague and we really would love to hear from you. The real benefit for us is this: as I say, this is an African solution.

Our goal is to be able to say we want to get into every single country in Africa and deploy DNS servers in-country. And in-country be able to serve at least every single other African ccTLD as well as stuff that we have on our name servers as well. So you can see the focus here. It's very simple: we want to improve the way African DNS works at no cost to the African environment.

So I like pictures and my little squiggle picture here shows the green countries are the ones that we've worked with and that we're successfully carrying in our Anycast platform. The yellow countries are the ones that we're working with. But the big question is why are there so many blank spaces?

And I'm hoping there are ccTLD operators in here that are in that map, in white. You can maybe tell me privately why you haven't responded to me or to my assistant nagging for the last year and a half. Please, we'd love to speak to you and if you want to tell be to bugger off and that you're not interested, I'd love to hear from you and get you involved and to see the benefit of the project – as I see it anyway.

So a little bit of information about it. Of course this is a little bit of a tech session. We have diverse, redundant nodes in Johannesburg and Cape Town. AfriNIC runs our infrastructure out of South Africa. We have upcoming installations in three more regions so if you know anything about Africa you know we have geo-political regions.

In the south we're going to be installing equipment that will service this kind of stuff out of the west. The north should be going live probably next week. We're working with people in the west and we're working with somebody in the east as well. And we have, for some measure of resiliency we have a hybrid OS application software strategy where we don't obviously put everything in one basket.

I say we're actively soliciting partners. If you're a regulator, if you're a Teleco operator, if you're an exchange point operator, if you're an African operator, really, and you see the benefit of having this in your country I would love to speak to you. We'd love to get equipment in your country so we can work on improving how stuff works.

What one of our typical installations, one of our nodes would look like… Out of the same Anycast pod we would provide you with the AfDSP services that I've just mentioned. We'd provide you with the 5855 services. We're also working with [RAC NCC? 01:41:27] to do measurement pods, so you'd be able to get the benefit of that. So like I said, that's a typical, install. I always get asked what we do for next-to-none profit, [laughter] a real non-profit.

We try to keep it simple. Puppet and VMware, which are tools that probably most of you are familiar with – Puppet to do the provisioning and VMware because, let's face it, it's easy and everyone's using it these days. Nagios for monitoring, DSC for the graphing – for the little orange and red pictures that you saw earlier on came out of the DSC.

Our software mix is typically FreeBSD and NSD because I like FreeBSD. But we have a growing number of Linux people coming on board so we also do Linux and BIND. Typically we split up Staff so there's always a mix and no single co-dependency. And of course we are guided by RFC2870, the excellent RFC that explains how a DNS should be run.

So effectively for me this is a three-stage project. The first one is obviously to get interested participants involved. As you can see we've gotten 21, so 21 of potentially 57 ccTLDs is good news for us. And that's excluding IDNs at this point, although we are able to support them and we have IDN of .tn on board at this stage.

So collect interested participants and get them involved and then get infrastructure in place. So we're currently at stage two of our project where we have deployed infrastructure in South Africa and we are deploying in the north in the next week or two and then in the west and in central Africa.

And the third part of it would be some sort of DNS portal, which would allow us to do automated interaction with participants so you'll be able to do things like change your DS keys, you'll be able to do things like update your zone or [force the fascias? 01:43:29] of your zone and stuff like that – what you would typically expect to have from any DNS provider. And that's coming a little later this year.

And that's it from me. So if you have questions, my name is Nishal. It's nishal@afrinic.net or our more generic, ticketed system is contact@afrinic.net. And we'd love to hear from you.

EBERHARD LISSE: Thank you very much. I have a question. So you're basically providing Anycast service for African ccTLDs? You want to provide that?

NISHAL GOBURDHAN: That's correct, yes. The third program, the AfDSP; that's exactly what it's doing. Free Anycast services for African ccTLDs. And I'll make this very clear – I should have done this early on – we have no interest in your registry business. What you do as a business is entirely your system.

We do not support you and I know you've heard about FRED, you've heard about CoCCATools – all of that – we don't get involved in that side. We're just supporting you in terms of helping to make your DNS infrastructure more resilient for Anycast DNS services.

EBERHARD LISSE: Thank you. Any questions from the floor? Nigel Roberts?

NIGEL ROBERTS: Sorry, I've got a problem with my voice this morning so please be patient if I break into a cough. Just followed on from something Eberhard said, DNS service for people in Africa is not just good service for the TLDs which are African ccTLDs – you have African ordinary users who will be trying to resolve New gTLDs or other ccTLDs.

Would you be interested in taking on Anycast broadcasting for other TLDs which aren't ostensibly African, but maybe are trying to be accessed by Africans in Africa?

NISHAL GOBURDHAN: Thank you, that's an excellent question and it's one that I often bang heads with some people in our organization about. Our immediate focus is African ccTLDs so that's where we are. On a personal level I definitely agree with you. I think it's important to be able to accept other ccTLDs and support them.

We believe we have the infrastructure and the expertise to do this but for now we're guided by what our Members want and the Board directive is to get this working primarily for Africans. But extending it is definitely on my list of things to do.

EBERHARD LISSE: I assume there will be a remote question now?

KRISTINA NORDSTROM: Yes, we have a question from the remote participants. It's from [Gandalf? 01:46:17] and he wants to know: "What is the smallest exchange point you are aware of in terms of traffic and peers?"

NISHAL GOBURDHAN: In Africa? I think the smallest Internet exchange point that I know of – and I'm looking to my colleague in the audience for guidance – is probably the [inaudible 01:46:35] Internet exchange points. There are two operators that are there at the moment. I do not think they publish statistics at this point so it would be a guess if I had to give you a number.

But in terms of participants it would be the [inaudible 01:46:46] Internet exchange point. The third participant that would warrant being an exchange point is the government and the university. I think the government is still trying to connect in but the university is there and is active at the moment.

EBERHARD LISSE: Any more questions? All right, thank you very much. I'm going to speak to you later, not on the microphone. Next would be a word from our sponsor. Where is Michele McCann? There you are. As you may already know, we are a little bit reluctant to allow too much marketing unless this speech is of course very riveting.

But we were approached by the ZACR that they had a sponsor and in particular, since this sponsor is offering services that none of us sitting here will actually take up there is no conflict of interest if we allow them to market a little bit to us.

MICHELE MCCANN: Thank you.

EBERHARD LISSE: One more thing is they have said that they have said that they have got data centers in Durban, in Johannesburg and in Cape Town. And they're willing to give us a guided tour. So if you have interest in that please remain here or maybe talk to her when we go and have our lunch in a few minutes. I see the box is being unpacked over there already.

MICHELE MCCANN: Hi, good afternoon everyone. To give you a very quick overview in terms of who Teraco is, we are a vendor-neutral co-location facility, as mentioned previously, located in Johannesburg, Durban and Cape Town. One of the key aspects around that, in terms of Africa, is there are actually no vendor-neutral facilities available.

So as you all talk about content and DNS services etc. in order to decide where you can actually co-locate those facilities, from a Teraco point of view. But going back to the presentation, which may be of interest to you, is we've done a lot of research around Africa and what are the business case opportunities in Africa.

So I wanted to take you through a few of those and we see it from a neutral perspective and maybe it might add some insight into that aspect. From an African point of view this is quite a famous map that we generally like to use. And some very interesting statistics around this is from an African perspective.

As you'll see, the size is that we are one third of the world's total land mass and over 1 billion in population, which equates to a total of 15.6% of the world's population. From this however there is only 7% of the world's Internet users based in Africa with an Internet penetration of 13%, however this is rising. So there's some interesting stats on that.

In terms of opportunities and what we see in terms of revenue generating – because everything at the end of the day does come down to money – is you'll see infrastructure projects are very key within our environment. We have a massive amount of financial growth or financial services growth and of that it's particularly mobile money, and Africa is quite a key environment.

And then obviously we have quite a big increase in terms of government spend, of which in South Africa, Teraco ourselves have seen quite a bit of that around the fiber infrastructure roll outs. Also another very famous map, which you'll see quite often. Generally it indicates the amount of capacity coming into Africa. Previously we… You'll see there is a very thin little line there and that was previously what Africa was running on and that is the SAT3 cable.

Now obviously infrastructures have boomed in terms of capacity, which then creates more opportunities for content providers and ISPs and DNS providers, etc., to be able to come into Africa. And with Teraco we've also seen a lot of international investment coming into our facilities and people co-locating their content and fiber infrastructure and the facilities.

Another very interesting slide – a lot of people say that there is a huge amount of cross-border issues and you can't generally service Africa from a central point. From a Teraco side of things we currently have… From Teraco you can pretty much connect to 59 different countries in order to distribute content and we've seen a lot of members come into the facility who are doing exactly that.

And the reason for that is that South Africa is a low-risk entry point; you know your hardware is going to be safe, no one's going to steal it, it's always going to have power as well as connectivity, as well as in terms of eyeballs, etc., that you can reach from a single facility.

Also a very interesting slide. In terms of our ICT landscape I think a gentleman earlier mentioned around this in terms of mobile users. Pretty much in terms of the Africa side of things, yes, it is predominantly mobile users and one of the interesting ones that I like is that in terms of the Facebook, which enjoys over 51 [sic] African users and is growing at 4.8% in terms of uptake.

And one of the things we've seen is the carriers actually starting to support this. A great example is that Orange have launched free access to Facebook. So a lot of carriers are starting to work with content players to actually start allowing the content to start distributing across Africa.

This is quite a famous slide, but one of the things we like to show is the difference in terms of being in a neutral facility versus being in a non-neutral facility. And a lot of you global guys will see it as standard practice to be in a neutral facility. In Africa it's quite a new thing. And what this has inherently done is slow down content distribution in Africa, purely because of the price points.

So if you see at the bottom price point there, in terms of interconnecting to each other with inter-facilities currently sitting at about $70 per 1MB. And if you start changing that into a neutral facility it's there that you start looking at proper Equinix or Telehouse business models, where you're looking at about $20 per private interconnection as well as zero cost to connect to any exchange point located in a facility.

And that's purely because as neutral facilities we like to promote content exchange points and as much interaction between each other so that the co-location space itself grows within the facility. This are some interesting stats from our perspective in terms of the amount of interconnect that are physically growing within the facility. And you'll see within the space of a year there is over 1,000 new interconnects and now growing at about 100 new interconnects a month.

So we see this from the physical side but what that generally means is that there are a lot of people coming into the facility and buying and selling content and product from each other, which is a positive thing from an African point of view.

Big thing: peering. Peering before we didn't… As Nishal and some of my colleagues have been promoting on their side is generally that peering is key from an African side of things and luckily we have a lot of support from the ISOC Members, etc., to start rolling out exchangers throughout Africa. This makes it incredibly important for Members to start distributing content.

So this slide is typically one of our clients and the amount of savings that they've actually achieved overnight, because now they are not stuck in a transit issue at the moment. So just to end off, because I don't want to hold everyone away from the lunches – the key facets in terms of growing an eco system is a neutral facility, low cost interconnect, open peering and then all of that will start driving content players, ISPs, Telco's carrier etc. to start putting infrastructure down in Africa and distributing out.

So as mentioned before we don't normally open up our facilities for everyone to see but there is my email address. If you guys want to have a day set up to tour either Jo'burg, Durban or Cape Town, wherever is relevant to you, please drop me a mail and then I can organize access and then you can have a tour of either one of our facilities. So thank you for your time and I hope you enjoy lunch.

EBERHARD LISSE: Thank you very much. [applause] Some stuff is quite relevant actually even though it is presented from a commercial perspective, but Internet exchanges often are only looked at from a political process or from an… [Let the ISP? 01:57:10] know, because they don't know really, but in the end it shows that it generates business and the promotion from there side, I think, is that if we foster cost-free interconnectivity we get more hosting out of it.

So it makes sense from a business perspective. But I have always been, as you may know, pushing for Internet exchangers because it's simply stupid to have to send an email to my neighbor in my street through London. From [Vintuck? 01:57:50] to Johannesburg to Amsterdam to London, and then goes over a different network to the different ISP who sits in the same data facility that my ISP does.

I find this extremely stupid. Yeah? I do not at all believe that this is an issue for legislation or for policy making. It should be purely driven by commercial things, because if you try to legislate it would create a whole host of other issues including bureaucracy and it won't really work. If we go slowly and make it attractive for the parties involved then eventually the [bin counters? 01:58:29] will figure it out.

And that said, we're going to have a boxed lunch now. It should be at 2 o'clock. Should the first presenter not have pitched up at that time we will rearrange the schedule a little bit and carry on, so I would appreciate it if you're more or less here at 2 o'clock.

[Tape change to ccNSO-tech-2-15jul13-en.mp3]

SPEAKER: This is the afternoon session for the Tech Day on July 14th at 1:53 pm.

EBERHARD LISSE: Okay, if we can also settle down now that we are all in our post-lunch lethargy. Sunday Folayan from Nigeria is... I don't know what he is but we emailed for 20 years more or less regularly and we've never seen each other before yesterday. We met on a development list; Africa development 20 years ago and we still communicate on the same thing. And when I asked for a presentation he was volunteered by his boss from .ng to give us a talk about ccTLDs and ENUM.

SUNDAY A. FOLAYAN: Thank you. Good afternoon everybody. I have 20 minutes to do this and I will very much try to keep it within that time. First of all let's go through the standard listeners' agreement and this agreement which you will all have to sign before I proceed. I'm sure you all agree that I'm not an expert on VoIP technology and that I am not pretending to be one.

And you all here sign and agree that I am a user who just got interested in the technology and its coolness and that you understand and agree that I am affiliated to my ccTLD and there it ends. No further issues. You also here agree that what I say may not be how it is but how I understand it. You all here now this afternoon agree that you will not believe what I say only but you will seek your own understandings and interpretations.

And finally you agree that you will challenge me if what I say is a lie or untrue in any form. So could you all append your signatures to the listeners' agreement? Okay. So having done that I can proceed that once upon a time camels were a means of transportation and in fact they are 4x4 all terrain general purpose vehicles. I don't think there is any disagreement with that. And many years ago everyone struggled to converge data in IP form into sound and move that over the public switch telephone network infrastructure used in those tiny devices we call modems.

But then VoIP came, which is the packetization and transport of public switched telephone system audio over well established IP networks and the audio and/or video streams that are recorded in digital format with possible compression and filtering before encapsulating it in IP for transportation over land or fabric or even the public infrastructure.

And I put up a diagram there which I call the VoIP matrix, which shows how you can move from IP to the standard public switched telephoning network, which we call the TDM technology and plain old telephone systems and various devices and functions for translating voice and data across those domains. For the technical buffs that makes a lot of sense but for you and I, who are simply users and don't understand much, just say, "Mmm, what a nice diagram," and let's move on.

So now what's happening is everyone is struggling to converge PSTN sound into data and move it over well-established IP links. And over the last 20, 10, 5 years, technology has just reversed itself. And so VoIP over IP has some advantages over the standard TDM that we have all known, which is the IP is scalable, it conserves capacity, it simplifies charging and billing – no longer do we pay by the minute or pay by the second but you can just have a flat rate.

Of course, while you can do video with IPs I don't know which Telco all over the world sells a telephone that is ready with video. You can do video calls with your standard telephone box or service. So Telco is still struggling with video. So all this provides a turf for value added service providers to sell softphones for PC to phone and PC-to-PC calls. And we've seen the success of companies like Skype really push telephoning to the next level.

Web-based applications for web-to-phone services. Clickatell is a very successful company out of South Africa that has done that very well and interconnection of office PBXs adds zero networking costs, so people can just dial from the Johannesburg office and they're talking to people in the Cape Town office at no extra cost.

And of course, it gives ubiquitous access to the PBX from home, travelling, people like me who just put on their headphones and can talk to any staff at their desk and still give the illusion that the boss is still within the country when in fact you may be on another continent. And of course, all those lovely features we love on PBXs like voicemail, call blocking, call forwarding and so on.

So that is a very exciting thing to do. The Telcos are [wise to eat? 00:17:13] and the major carriers and what they do is they take their traffic, change it to IP, stuff it into the Internet and at the other end the destination carrier retrieves it from IP, converts it back to TDM and it looks as if two telephone companies just exchanged traffic, when in actual fact they sent it over the Internet.

Of course, as the big boys play, little kids also play. So you find in Africa things like call stuffing is very popular, where people set very tiny boxes in their bedrooms and take advantage of the huge rates between international calls and local calls and they carry huge traffic over IP networks and reinsert it into the networks of their home country.

And this is something so many countries are fighting, so many countries are struggling, and I think it's just a lost battle because you really can't fight technology. When the time comes people use technology and you either adapt or you die like the camel as a means of transportation.

So eventually what we're going to see is what is true universal access, which is people being able to communicate, whether it's a mobile phone or a desk phone or PBX or an IP phone – it doesn't matter where you are, you can see it and take your call. You are not paying for roaming, which is some rip-off that some of us have been complaining about for the last few days – but then that's what the rules are now and that's how the game is played.

So enter ENUM, which is an IEFT standard actually defined by RFC 2916, which is for mapping public telephone number address space into the domain name systems. One of the things of talking at a meeting like this is you don't really know whether you are carrying a call to Newcastle, whether you are saying something people already know or things people don't know at all.

So it's really a tough nut to try to know the depth and scale that you can go technical. And when you are warned that you are the first speaker after lunch you don't want to send people to sleep so it has to be a very light topic. So the role of ccTLDs that the opportunity that this technology is presenting now is that ccTLDs are in a very, very unique position of mapping telephone numbers within the country onto the Internet using the ENUM technology. Then people can take advantage of broadband Internet access and be able to make very affordable calls whilst still retaining their well-known telephone numbers.

So people don't have to think twice or crack their brain to know the details of whom they want to call. And what this does is allow the conventional telephones to call IP terminals like PCs and vice-versa. And should telephone numbers be used in this way it would be possible to offer a centralized database of numbers and end points.

Of course, as far as I'm concerned who should exercise control over telephone numbers used this way. Of course I also would state that it should very much be an opt-in, opt-out system. If I don't want to be reached by ENUM it shouldn't be compulsory. And so there are other territories and colonies that I believe the adoption of these types of technologies on a larger scale within the continent and indeed the world will offer.

It's very prevalent these days to walk out and see so many Wi-Fi networks. Some open, some locked. But it would be great to have a universal standard where people can set up their Wi-Fi systems and be able to offer secondary servers that passers-by with Wi-Fi network phones can actually authenticate of those microcells and be able to connect to a Wi-Fi network as they move on and have access and be called over Wi-Fi.

Of course, what that means is that we would be seeing some phones appear in the market in one or two years that are labeled with “built for ENUM”, which means those phones are ENUM ready, you just configure your ENUM server and the phones handle all the other things and they are able to communicate with the end point and allow you to talk while moving.

It would also not be too difficult to see applications that would run on smartphones able to offer this type of service, which means finally we can have true global roaming at local rates. That in a nutshell is what I have for you after lunch and I hope it didn’t send you to sleep. Thank you.

EBERHARD LISSE: Thank you very much. [applause] Any questions?

DAN NEWARK: Sunday, my name is [Dan Newark? 00:22:50]. I work for Internet [side of it? 00:22:48]. I’m here as an individual because I’ve worked in VoIP for ten years. What you’re describing is definitely the vision that many folks had for ENUM.

The challenge that came up in the IEFT Working Groups where this was being debated was that the public ENUM, the use like what you described here, ran into the issue that it could be easily used by spammers, by people who want to do telemarketing, by people who want to do that. There are even some applications out there that will walk the ENUM tree through DNS and pull down all the numbers and then start calling out to all of those numbers throughout there.

So some of this original vision that people had for public ENUM like this is not happening right now because of this large issue around this. Now, ENUM is being used in a huge way by carriers and folks internally to do this. So I think some of that is not happening in the way that certainly many of us hoped.

However, you hit a key point in there, which is that ccTLDs are poised in a good place to be an identity provider for telephone numbers… Because you are a logical place for that with what's happening in there and I would encourage you to look… There's a new Working Group happening in the IETF called STIR. It's about looking at secure telephone numbers and how we can secure telephone identities.

And many of the people involved with ENUM are involved with working on that and looking at how we can create a way to have secure telephone numbers – caller ID, if you will – for this, which goes back to the issue around how you have identity and link it to telephone numbers. And I think there is a great role that ccTLDs can play in that space.

So that's where I'll end up with that. There is some good work happening in that space and I'd be glad to talk to you some more about it at some point too.

SUNDAY A. FOLAYAN: Thank you very much. I'm glad you agree with me on that and that ccTLDs have a great opportunity to really, really show some direction there. Thank you.

RICHARD LAMB: Rick Lamb, ICANN. Just as myself. Cool. Do the access points have to get modified to support this? So you're envisioning a world where I can walk by my neighbor's house and borrow a bit of bandwidth from that, right? What has to happen at the CPE side, at the access points? What has to be modified in order to support that?

Authentication has to happen seamlessly across the Wi-Fi cells, for example, right? I'm trying to understand how… I like it, but how do I make it happen? I was trying to figure out what the limitations are.

SUNDAY A. FOLAYAN: First of all, where there is a will there is a way. I'm happy that the idea makes sense. The mobile phones don't do anything much more technically challenging; they actually hand you off from cell to cell. You had the idea that you're moving on continually seamlessly. So I don't think it would be a challenge to have a system where you end a connection from one Wi-Fi cell to another.

Obviously there is going to be a common SSID, for example, and a common means of making sure that there is the handing off and the reconnection. That's not too difficult for some smart folks to walk on and I want to even bet a penny with you now that one or two Google engineers have already done that.

EBERHARD LISSE: Where's Warren? Warren is one of those Google engineers but he's not in here.

PATRIK FALSTROM: Patrik Falstrom, [Netnode? 00:27:00]. I'm one of the persons to blame for this ENUM thing because I wrote the standard. So let me clarify a little bit about what Dan said. I think you're doing really good things and you should just move ahead exactly like you are doing, because many of the telephone problems and spam issues that Dan mentioned are not so much related to ENUM but it's a generic voice or IP [CIP? 00:27:26]-related issues.

And people are solving that quite a lot today by ensuring that two peers that do exchange, for example telephone companies or something, that you exchange the traffic and also do not expose the [ingras-to-ingras? 00:27:40] point for the telephone calls between themselves. So you use a different [ingras? 00:27:45] point for your voice or IP calls for public calls from the ones that you [deal? 00:27:50] do with the ones that you actually have agreements with.

On top of that the IETF is launching that Working Group that Dan was talking about. The biggest resistance about ENUM in, for example Sweden where I'm coming from and other parts of the western world, it's that you have old incumbent telephone companies that don't dare to give up the responsibility of telling where phone calls or any service related to the phone number should be routed to the individual that would actually like to control its phone number.

And specifically they also don't want to do that with enterprises. That's a resistance that I specifically see in for example Sweden and the US and other countries which to some degree dominate the IETF. So I see much more movement regarding using ENUM for a number of portability applications that you talk about in non-western countries.

So move ahead. Continue.

SUNDAY A. FOLAYAN: Thank you very much for that comment and I think I have very nice news out of Nigeria. Just two months ago Nigeria liberalized that and a number of portabilities now stand there in the country. Once you have a telephone number you can move it from one operator to the other without changing numbers. You cannot do that in Nigeria, which means that the numbers are attached to the individuals.

This again makes it very easy to revive ENUM, domicile it with the database and be able to offer that type of service. And I think as other African countries overcome that barrier of allowing number portability, these things would come back to the front bonus and it will be possible to do it on a global scale. Thank you. [applause]

EBERHARD LISSE: Thank you very much.

ROD RASMUSSEN: Hello everyone. Thanks for having me back at Tech Day. It's good to see a lot of friends and hopefully soon-to-be friends in the audience. I have a couple of topics that we'll try and cover in 20 minutes. So let's get at it. First thing I'm going to talk about are some thoughts on mitigating botnets within TLDs and various issues around command and control servers and the like.

Then I'm going to give you a very quick overview of what we've come up with in the Expert Working Group on next generation directory services. And we're having the public forum meeting at the same time as this so you'll get a quick view and hopefully be able to throw a couple of questions at me – I'm one of the Members of the Expert Working Group.

So thinking around botnets and a TLD. So what's the problem? As many of you know, people set up botnets and typically use domain names to manage them. To manage them you set up what we call rendezvous locations where you go to find out where to get instructions or download updates and things like that, and various other communication functions that botnets do, which is often relaying instructions from one server control and command server to another server.

There are a few ways that operators of botnets have to run their systems. The easiest and most basic way you see is they hardcode specific domain names into the code that's running on the bot and that reaches out to specific servers. This is especially true in today's model where you have resale.

So somebody writes the code and then the resell if over and over again to different criminals. We call them “script kiddies” or various other criminals that aren’t as sophisticated technically but they can buy this wholesale from somebody and then get into it, hardcode various domain names to use as part of the infrastructure.

You can also use a rendezvous server to update those domains over time and then we have what we call a domain generation algorithm or DGA and I’m sure most of you are familiar with Conficker – back in the day that was probably the most famous example of a DGA, where they had thousands of domain names being generated on a daily basis and the idea there is that you use…

The software that’s running on the bot runs a code based on the day or something about the operating system or something else – we’ve actually seen someone use Twitter feeds to generate domain names on which to rendezvous on.

Those are particularly fun because they do a top-trends that are trending and they create a domain name based on that and then try and rendezvous with it so you actually play a racing game between the good and the bad guys to try and register that domain, so that’s a fun one.

But that’s a pretty hard one to manage from the botnet herder perspective so they have a DGA when they generate that. And they will often generate this string within a TLD. One of the good things to know is that these can be very easy to find. You get the code and reverse it then you know what they are, and that’s the typical way of doing it.

But if you don't have access to that you can take a look at domain names that are being registered and they often are very easy to spot because they don't make any sense at all. They are a series of letters and numbers for example and sometimes they fall within discernable patterns where you have a serial number that's rotating, things like that.

Various algorithms or machine learning can be applied towards your zone file for example and these will pop out really easily. Oftentimes they use the same name servers over and over again, despite the fact they're generating new domains all the time but they put them on the same infrastructure. And if it's not the same name servers they will usually ask for flux or double-flux hosting.

So these are all things that within the running of a TLD you can actually see these patterns showing up in your zone files as registrations come in for those domains. And you also have some more tools at your disposal. If you're taking a look at DNS queries against your zone. So oftentimes DGAs will generate queries for domains that don't exist because they have this… Conficker is a great example; thousands and thousands of them.

You will see those resolutions coming into the TLD as requests so at that point you can even anticipate if you're watching your resolutions when a new botnet has taken up, because you'll see these non-existent domains showing up. So if you have real time access you can actually take a look at that. And another thing that a lot of people will track as well is looking at particular registrars where these show up.

So there are various things you can put in place to watch for this kind of activity, right from your own view. And the final thing is that people will tell you about then. Lot of people who are trying to shut them down, obviously. One of the things I want to emphasize though is this notion of shutting down a botnet via this, by taking out the domain name for it.

A lot of people though are trying to move towards sinkholing. The idea behind sinkholing is we set up a server on that domain that's being used by the botnet and when we get requests from the bots you track that and you notify the networks where that request cam from that they have an infection and here's the timestamp and the IP address that had attempted to reach out to you.

Now, obviously policy within your operation and your country will vary as to whether or not you can do this, but this is a really valuable activity to be able to do but you have to have policy around that. It's not something you can do lightly; you have to have a really good idea of how this will work, how to make it equitable and to make sure the bad guys don't show up at your door asking to sinkhole things from other bad guys.

There are also people telling me about these things. There are lots of resources out there. If you want to find out about command and control domains that are on your network there are a huge bunch of free services out there and I encourage you to work with some of these folks. The people at SURBL will be very happy to set up a feed for you, command and control domains within your TLD.

I've talked to them several times about that. It's a service they offer. I've listed several others here where you can actually get a feed of data or if you suspect that something may be a sinkhole you can actually… For example Google Safe Browsing I believe allows you to query against another domain name and they give you a response back as to what they think about that. And Microsoft have a similar program that they've set up.

There are commercial operations out there as well. These are paid-for services that will give you real-time updates on things. And there are lots of these as well. I left Netcraft off there, I forgot them. They do this as well and I'm sure lots of you saw the Architellos report. My own company does it. Symantec, Websense… There are lots of people out there that do this that will provide you that kind of information.

So there's resources, there are paid-for services. So if you're trying to take a look at cleaning up and keeping botnets off your TLD there are lots of different options you have, from running it yourself to actually getting people to tell you about them.

And the last thing I wanted to talk about here is what's your policy on this. We have different policies out there from "I won't touch this at all" to sinkholing when someone complains about it. And the other question people have is: "How do I know this is really a c2? Somebody tells me this is a command and control server but how do I confirm that before I do something about it?"

And also: "Who do I have do something to it? The registrar? The registry?" It depends on your model. But you need to be able to confirm this thing, typically. So you need to have people on Staff be integrated with a cert; BrazilNIC is a good example, or you get some kind of outsource threat intelligence where you can say: "Take care of this. Tell me whether this is real or not."

So again this depends on resources that you have. And do you suspend it, delete it, sinkhole it or even transfer it? We know cases where different TLDs, different registrars will transfer domain names that are registered by criminals for running botnets to security companies or to law enforcement, things like that.

You should have policy around this so you can answer consistently when you get requests. But you should develop policy if you haven't already. Okay. Very quickly, before I go onto the next section, were there any quick questions on the first bit?

Cool. We'll move onto the Expert Working Group and this is a very brief version of the presentation we're about to do in 15 minutes in the other room. I'm sure many of you heard about this. It was set up by the Board to take a look at a replacement for WHOIS. Think of it as a clean slate. We're going to burn it to the ground and start over.

How would we do domain registration, directory services, if we were starting today and not in the 1970s when we were just universities talking to each other and trying to fix things? So we really took a look at redefining all the purposes, redefining the system based on all the purposes and the uses, etc., that are out there.

And we formed this Working Group. The initial report was published on the 24th and has been out there for a bit. There has been a lot of comments and we want to keep them coming in. I do want to say that a lot of the ideas we went with were inspired by what various ccTLDs are doing within their own registration policies and policies around display capture, policy around different types of users out there and etc. for developing what we've come up with.

But what we did was really start with taking a look at it as, if we were to provide a brand new service in a market place that was being served very poorly – which I think describes the situation with today's WHOIS pretty well – what would we do as an organization where we're building a company.

So we did a used-case analysis. We started with who uses domain names? Who interacts with domain names? Who needs to get information around those domain names? What are their purposes for doing so? And this included everything ranging from: "I want to register a domain name to run a website or use it for my email," and all the things you need to do to do that, to: "I'm a corporation, I want to manage my trademarks," "I'm a person who wants to trade in domain names," "I am a criminal who wants to spam you."

So we went through a whole series of use cases and found out all the various things that people do and then came up with a purpose driven model. So in the end what we've come up with is that we're going to abandon this "one size fits all" system where everybody gets the same access, it's all anonymous and it has to be supported by this wide range of registrars. And we'll replace it with a system where you have accountability on all peoples' parts: the people that are registering it, the people who are maintaining the data and the people who are requesting information about domain names, which we don't have today.

And to allow for this to provide better privacy we've looked a lot at privacy concerns, commissioners and different privacy concerns around the world. I know that's a big topic that's starting to rear itself around here with the 29 folks coming out of Europe saying you need to be looking at how you're doing this.

So we're really trying to build or create a system that would allow for that and increase the accuracy. Because today's system's biggest problem is that it's very inaccurate based on lots of different reasons. So we have…

What we've done is create this framework, and I'll talk about them all in a minute but really what we should be concentrating on with feedback is looking at the framework we're trying to do here and base it on uses user purposes, etc., and how you provide access to that with accountability.

We have all of these different consensus and policies and principles that we came up with within the very diverse group, representing pretty much all the key interests in the space. But how do we provide for validation and accuracy in the data? How do we allow different people with different kinds of needs to access the data?

For example if I'm just a consumer and I want to find out a little bit about a domain name that's tied to a website where I'm trying to transact business, I can go and look anonymously and get a certain amount of data. And by the way, there was a misconception – we're not proposing removing anonymous access to data; we're saying keep that in. Limit the amount you might be able to see anonymously.

If you want to see more data for different purposes, say you're filing a UDRP claim or you have some abuse issues you're dealing with, you authenticate yourself to the system and it gets more data about who you may be able to contact to deal with whatever the issue is. If you're law enforcement you might be able to get even more data by being an accredited law enforcement officer, via a system that we've come up with.

So those are the main principles and those are the things we want to accomplish. We also came up with a model that would support this, called the aggregated registration data system, or ARDS. But the idea here is that registrants still work with registrars to enter data and collect it up and make changes and all that stuff. And it's stored at the registries at the authoritative source for all data. And that's pretty much what the THICK WHOIS model is.

So that's where we are today. What we're really adding is the idea of being able to provide purpose driven data disclosure or access to the data through this ARDS system where you have copies of the information from all the registries and you can put it all in one place and be able to pull from that and do all kinds of efficiencies and it allows for authentication, accountability, etc., on behalf of the people accessing the system.

And there are a lot of benefits. Read the paper for what else this brings. There are some potential drawbacks as well and we point those out. There is one big data system there and that creates some security issues that you want to think about. So we have this consensus view from the group. We've got the report out there that reflects a lot of compromises. We've had a lot of discussion.

We had a whole day yesterday. We're still working on this and we're working more and more with taking input we've got already. We had a full day yesterday and we've got several hard problems that we're dealing with and still have to deal with. And we know it's not perfect and we know there are issues and everybody's going to have some complaint about something – that's the nature of the beast when you're trying to do something this big.

What we really want though is your constructive feedback. Not, "This sucks, make it go away," or, "This is the best thing ever, we love it!" That's great to hear but that's not really constructive. We want to know why it's good or we want to know why it's bad. But in general what we need is feedback on how we can make it better and what did we miss.

We know we didn't think of everything. We think we're the experts but there are 13, 14 people with a lot of Staff help. And there are a lot of issues out there and you don't think of everything. So we really need that feedback. August 12th is the current deadline on that. We've got addresses there. We're going to continue work. We've got this open meeting here in just a few minutes. We're shooting to deliver a final report in or before October so that we have enough time to get feedback and present this again at the Buenos Aires meeting.

And this gets delivered to the Board. This is not a standard gNSO thing, this is actually commissioned by the Board and it gets delivered to the Board. From there, if these things are accepted or what have you the Board will send it over to the gNSO to PDPs around the various policy implications that will come from that.

So this is not a "we're going to change WHOIS by next January" but this is hopefully the beginning of a process to come up with a good model that will be used. I have a lot of potential discussion questions but I see I am running out of time here. I will take your questions and I've given you the presentation; feel free to distribute it to the members here.

One or two quick questions because I don't want to take the next person's time either.

EBERHARD LISSE: The presentation is of course on the website so if you go to the schedule you can get it. And one thing is clear; this is not a ccNSO PDP so the ccTLDs will not in any way be bound by that. However, if you use standard tools like FRED or CoCCATools they will implement it, I'm sure. I'm quite sure that [Das Miller? 00:49:58] and his team will implement it in CoCCATools so if you use those tools it will be done.

Why not use it if it's available? I don't care if it's WHOIS or whatever it's called, as long as it works. Any questions?

ROD RASMUSSEN: I would encourage ccTLD operators to take a look at what we come out with. Depending on the policy in your territory you may want to take advantage of a system if it gets stored up at some point.

SPEAKER: Hi Rod. Is there a timeline for this?

ROD RASMUSSEN: August 12th for feedback – we're getting a lot but we want more. Then we're going to have a final report out October timeframe, hopefully, in order to fork that out and present it in Buenos Aires. Then it goes to the Board and from there we don't know what the timing is going to be but given the amount of interest in finally solving this problem I think we're going to turn that around to the gNSO fairly quickly and hopefully have a process that's fairly straightforward.

SPEAKER: I'll put it another way. I've seen this morning the mention of a five-year plan or a three-year plan and stuff like that. Do you know if that is part of a three-year plan or a five-year plan that ICANN has?

ROD RASMUSSEN: I'm not sure if it's in the five-year plan. I hope it's less than five years but we have been dealing with this problem for a lot longer than that too so. Okay, well, I've got to run. Thank you for your time, I really appreciate it.

EBERHARD LISSE: We have a remote question so I think we should take it.

KRISTINA NORDSTROM: A quick question from Antoinne: what jurisdiction would this ARDS be in?

ROD RASMUSSEN: That is open. We'd recommend an open international jurisdiction. I know there is a lot of sensitivity to the US based things at the moment so we are definitely saying this should be some sort of international… It has to be somewhere. Maybe Switzerland or somewhere. [laughter] We are definitely not prescribing that it would be an ICANN thing run out of Los Angeles. It should be something independent from that.

EBERHARD LISSE: All right. Thank you very much. The next presentation will be about hijacking by Marten van Horenbeeck.

MARTEN VAN HORENBEECK: First of all, thank you very much everybody for inviting me to speak here today on a topic that's actually fairly close to me and my company's heart; DNS hijackings. I don't know how many of you were at the Costa Rica ICANN where my colleague, Morgan, presented on this a couple of years ago, I think in 2012. But what I'll try to do today is give you a little bit of an update on what Google has seen in the field of DNS hijackings and share with you some of the things that we learnt from registries.

Now, before I get started I do want to make clear that it is not my intent at all to point fingers or to identify particular places where things went wrong, though there will be a slide with an overview of recent incidents. But at the same time this is an issue that affects both you as registry operators and Google as a company that clearly has a lot of domain names in various ccTLDs.

So what we really want to do is work with you to find solutions for these types of problems. And I'm really here to meet with you throughout the next week so I'd say if afterwards you have questions or comments, please send me an email and reach out to me and I'd be very happy to engage in productive discussion in finding solutions to this.

First of all a quick introduction. I work in the Google Information Security Team and really what I do there is I work with a smaller group of that team that is responsible for making sure that when a user visits a Google website that they actually end up with Google. So that involves two things: it involves SSL and certificate authorities and on the second and it also involves domains and how we protect them at Google.

Outside of that I'm also the Chair of the Forum for Incident Response and Security Teams – FIRST – which is an organization which brings together security teams from various countries' enterprises and even a couple of registries. Maybe to give you a little bit of background as to how large this problem is slowly becoming: this is an overview which is non-exhaustive of a few incidents that happened from 2010 to 2013.

And these are really just the ones that are notable; where we actually noticed very clearly that our Google domain ended up being hijacked in this particular country. Some of these are not necessarily related to the registry. Some of these include registrar. Hijacks are also a little bit more rare. But as you can see there is clearly an increase in the amount of incidents.

What we're also seeing is that there is an increase in the amount of people actually looking to compromise a registry and taking over particular websites. Now, the impact of these incidents really differs from incident to incident. And in general we've been quite lucky that most of them have been clear defacements, in the sense that an attacker compromised a registry, took over a select set of websites, pointed it to a server that was under their control and used it to display a defacement page.

In many ways that's actually the best possible outcome because in that case the only thing that is really at risk is brand image and the users that are somewhat concerned about the fact that they're browsing a particular site and they actually end up on a site where there is a political message.

But in many ways this is preferable to the two other things that can actually happen. One of them is that the server that the user ends up with could actually steal authentication tokens. Now, in many cases SSL can help protect against this but not all popular websites necessarily use

SSL by default. So that means that there is a certain risk that tokens may actually be stolen and reused by the attackers.

In addition, an attacker could also deploy an exploit kit on the website to exploit code, to exploit vulnerability in the browser that the user is using to visit this particular website. And that could be used to create sizable botnets. Luckily, we haven't really seen the second and third being very common yet but then again the impact of a hijack is really difficult to assess because something that looks like defacement may actually still be collecting authentication tokens.

And also, defacement may look different in different parts of the world because the attacker who owns the final server where the user ends up is really the organization or the person that controls what will happen to the client when they visit the website. So it's not always very clear what is exactly happening.

Now, this is just one example. It was in November last year where we actually had our Google domain in Pakistan hijacked and it took us several hours to really get the necessary people on the phone and get everything addressed completely. This is another example of last year where both Google and several other companies' domain names were hijacked after an incident at IEDR.

And as you can see, these incidents affect both the registry and the organization like Google at the same time. So that's why we really want to work with you and find solutions for this. One of the things that we've been doing recently is we've been trying to reach out to every single registry that had a compromise or had a certain incident. We've been trying to collect some data on what seems to be happening regularly in these hijacks.

And these are just some of the findings that we had from the last year or so. First of all we noticed that a lot of these hijacks actually happened at the end of the business day where the registry operates. For instance quite often we see this happening at 6:00 pm when the registry management has just left the office and it's really difficult for us to get someone on the phone until the morning after when they start working again.

And of course this is problematic because that means that during the night there is very little that we can actually do to protect our users that are ending up on a site that we no longer control. Partially because of this the average time to address an issue is often quite long. I say here it's often six plus hours because that's an average, but there are many sites or many registries where recovering actually took quite a little bit longer.

Also, registries for very reasonable reasons want to prevent additional domain names from being hijacked. And what you need to do there is very different than what you need to do to recover some of the domain names that are hijacked. For instance the first step will be to take the website offline that actually allows the registrars or end users to change their domain name registration.

But at the same time, that also prevents companies like Google or a registrar from recovering some of the data that was actually changed during the hijack. Also restoring the previous things is often errorprone, in the sense that sometimes it's a back-up, it's not always clear when those registry files that are being copied over and recovered were dated, so there is a lot of confusion that sometimes ensues.

And we've seen in at least two cases in the last year that during the recovery something actually went wrong and incorrect data files were restored, causing the domain to still be unavailable for a different reason. And very important for us is that even when the domain name settings are recovered, very often clients and intermediate service providers like ISPs actually cache the poisoned results for up to 24 hours.

And in many ways the hijacker has control over this because he can change some of these time-to-lift settings and actually influence how long it takes to recover. And this is a big issue for us because in many ways we can reach out to certain ISPs and ask them to flush their cache settings, for instance.

But when it's a problem it's sometimes a bit more difficult for us to do this. And we do sometimes try to get the help from the registrants to see who we should be reaching out to, and if they can help us reach out to some of them to actually flush the cache so the domain name becomes available again.

And finally for us it's not always clear when the incidence is really contained. We've actually had – and you may have seen this from the second slide – situations where one registry was hijacked more than once and also there are situations where a registry was hijacked and suddenly other hijackers see that something happened.

Then as soon as the registry comes online again they start probing it for software vulnerabilities in the website because they know that there was one so not all of them might be fixed by the time the website comes online again. So there is a lot of complexity with actually recovering from these attacks and so I believe that one of the better things we can do is try to understand which are the most common types of issues that are being used to exploit us and how we actually address them.

So in the calls that we had with many of these registrars and registries we actually asked them how their site got compromised and what actually happened. And we found out that in most cases these were vulnerabilities, not so much in the DNS services that they were operating but vulnerabilities in the websites that they actually manage.

And the number one vulnerability is something that's fairly common and fairly well documented so there is a lot of information on this on the Internet, but it's issues with the authorization schemer. So what we've seen quite a bit is that an attacker creates a second account or his own account on a particular registry, logs in and then once the attacker is logged in he manages to play around with some of the ID settings in the CDP request header.

And by doing so he manages to gain access to another account, for instance the account of a registrar that manages a popular domain. And once he has access to that account he can of course modify some of the domains that were registered by that particular registrar on the registry. So mitigations for that really include fairly basic things like configuring appropriate roles and privileges for each user account that is going to access the web portal and second, making sure that that user account and authentication mechanism is enforced every time a new resource is being accessed.

For instance when a user logs in and creates a domain name, that user should not be able to access other domain names by changing settings on the client side or changing parts of the request string. And this, even though it seems fairly simple, is actually an issue that we've seen a lot and it mostly seems to occur with registrants or registries that have developed their own software and have not really been maintaining it a lot in the last few years but have been operating it on a day-to-day basis.

Now, for this particular issue, one thing that I would like to recommend is there is a great organization called OWASP – the Organization for Web Applications Security or the Open Web Application Security Project – and they actually have a really good testing guide with a lot of information on how to find some of these issues and how to address them in your code.

The second most popular issue is fairly well known and I've you've dabbled in security you will be quite familiar with it: it's SQL injection. This is a situation where certain parameters on the website are not necessarily being checked correctly to see if they contain certain characters like ticks that might allow a user to escape a parameter in a SQL query, then replace it with a change to the SQL query.

So here on the slide you can see the first line, which is what a query would look like in web server logs or what it looks like towards the web server. And one then ends up being translated into these SQL statements. There are many ways to deal with this. One of the common ones in the past has been to sanitize information to make sure that this type of tick actually gets filtered out.

But more comprehensive is really to use prepared statements or stored procedures in SQL – procedures that you essentially prepare ahead of time and you only have to make small pre-approved modifications based on the entry that is made on the website and then execute it. Sometimes a very quick way of starting to protect against these issues is by deploying a web application firewall.

And we've seen a lot of registries that were in fact compromised do this, so we started using tools like mod_security, which is a patching module that allows you to filter certain things from the request string. And we've seen that used with fairly great success at helping protect against these issues. But most important really is to know where these vulnerabilities might exist and then address them.

Finally, the third issue is in many ways the least technical but it also requires some work to address. So if the account of a particular registrar on the registry was hijacked for instance by password brute forcing or alternatively when there was another registry that was hijacked the password database was copied over and posted somewhere and that password was used on other registries as well.

These are fairly common things and there are actually quite a few fairly easy ways to deal with this problem. The first of course is to properly hash any passwords so that when someone breaks into the authentication database they cannot just copy the clear text passwords, or even hashed passwords, if they weren't salted – that means adding a particular value to it so that it becomes very difficult to use pre-generated MD5 or hash cookbook of a set of passwords to try and decrypt what password a particular registry used.

Second, using two-factor authentication is very helpful because that essentially moves the password as a single point of failure. Using IP address restrictions for certain high-level registrars… So if you have a registrar that owns a number of domains or has registered a number of domains on your registry that are very high traffic, you might want to consider working with them in putting in place restrictions on the IP address they can connect from.

Locking out accounts after they unsuccessfully try to authenticate: if a user has tried to log in five or ten times with a wrong password there is an indication that something might be wrong so you might want to temporarily lock out that account and see what happened before you unlock it. And finally implementing something as simple as password strength requirements.

So overall a lot of these things are really software vulnerabilities in the website and what we've seen quite a bit is that this happens on sites that were built quite a while ago or even fairly recently but don't have fulltime developers working on fixing bugs and identifying bugs proactively.

So there are a couple of recommendations we would like to make to registries that might be interesting to consider. And the first one is to deploy registry lock. The concept behind this is that certain domains would no longer be able to be changed interactively on the website without some other type of modification taking place or authorization taking place for instance by sending a signed document or a fax from the registrant to the registry.

This takes a little bit more effort but in reality most of these domains are very unlikely to change so it does actually provide a lot of protection for high-level domains such as ours and some other online service companies. Verisign actually has a great description dating back to 2009 of how they implemented it. It's posted to the ICANN website. So you might want to review that if this is something of interest.

Second, having emergency contact information available to your registrars and high traffic domain name owners might be something of interest as well. For us this really helps because as I said, one of the most common things we see is that the mains are actually hijacked in the early evening. And in many ways, if we could reach out to the registry at that point in time we may be able to help you get things fixed very quickly. But this isn't always possible.

Quite often the contact information we use is the information from the public website. So if you have an on-call number, something of that nature that you can provide to your registrars, that's something that's very helpful. And I do recognize that not everybody has the resources to have people on call, but by sharing it with a select set of organizations that won't call you for the least small problem that might still be useful.

Second – sending notification emails when changes to domain names happen is very useful. We actually monitor [Darkskin? 01:09:48] constantly and we have been able to see that a change was happening to a domain name, contact our registrar and contact the registry and actually be able to prevent the hijack from actually taking effect into the DNS system before it actually is replicated.

So we were able to protect all of those domain names that were operating on that registry, purely because we got an email saying something was wrong. And then finally, think about using high traffic domain names as useful canaries. In general you will rarely, if ever, see Google domain name settings change.

So if there is a change and suddenly you see a DNS system being configured that does not look at all like a Google DNS server, then that's a pretty good indication that something is wrong and I'm very happy to give you my cell phone number if you ever see this and you want to double-check with us, because quite often this will be a very good indication that things are not the way that they appear to be.

Last – this is probably the most difficult because it does take some amount of effort but it's really quite important – is to develop a security management framework for your registry. And that doesn't necessarily need to be according to the ISOC standards, though that is of course a very rigid framework. Quite often it can just mean that you sit together with your engineers and you talk about what are really the security issues that we have.

For instance, is the software that we're using today still supported or if something would happen with it do we need to spend a lot of effort ramping up on fixing vulnerabilities in it, and other software projects that we can use that are open source and that we can modify for our particular registry project. Second, one of the biggest indicators that we've seen for registries not getting hijacked again was actually when they consulted someone external for security testing and actually had some security people validate the website.

In almost all of the cases where the registry did this we did not see any further problems with the registry. So this is a very quick-win type of approach to solving a lot of these issues. Finally, it's important to have a good understanding of the weaknesses of the infrastructure that you operate. So what really are the weak points?

Is it a website or are you running an older version of DNS server? Or is it perhaps that insiders within your company have almost unfettered access to making changes and thereby exposing them to things like spear-phishing or cross-site scripting attacks in emails that are sent to them? Those are really good things to understand and once you have a good understanding of what the weak spots are you also have a good idea of where to invest and what to build out.

For instance, implementing a web application firewall in front of the website to mitigate current issues while you consider moving to an application or developing an application that's a little bit more modern than an older software package that you might be using. And finally, having a plan in place is really important because I've spoke to many registries in the last year that were not expecting to be hijacked because they actually were running up-to-date software, but the attacker there was particularly interested in getting access because it was for instance a very large registry.

So having an incident response plan on how you're going to communicate with your registrars, how you're going to enable registrars to recover some of these high traffic domains is very important. And one of the things that I have been trying to do with some of the registries is to push them to what's reaching out to the local computer security incident response team; C-CIRT or CIRT.

Most countries today have one and really most countries only have one registry. So from that perspective there is often a lot of potential for partnership between this one organization that's very focused on security and this one organization, the registry, that's actually running very crucial national infrastructure when it comes to the Internet. So reaching out to them might be a very good way of getting support and a good up-to-date awareness of what is happening in the security space.

And that is actually all that I have for you but I really want this to be a kick-off conversation with you to see where we, as Google, can help support you in solving some of the issues that you might be concerned about and working together with you to make the Internet a safer place for everyone. So my email address is on the slide. Please feel free to email me.

I'll also be here the rest of the week and really I'm here to talk to you so feel free to walk up to me and ask questions or ask questions right now if we still have a few minutes.

EBERHARD LISSE: Thank you very much. I will use the prerogative of the Chair for two things. We had this Microsoft initiative present in Beijing and as far as CoCCATools is concerned they have progressed very far. We have found a few updates and they're very close to finishing this. So if one uses the latest version of CoCCATools these vulnerabilities will have been assessed, if not closed.

The other thing is we occasionally get emails from your registrars for requests and then I find out I've never communicated with this person before and there is no PGP or other signature on the email. How do I respond to that?

MARTEN VAN HORENBEECK: You mean you receive emails from registrars that are not PGP-signed or…?

EBERHARD LISSE: No. I have received an email from a person at the registrar that was responsible for your domain name and I didn't know the person. And they asked something not unreasonable but they asked something to make some changes and they were unable to provide me with a signature. Oh, they're standing there. Yeah. Maybe he should answer himself but the point is that this is something that we encountered and you should perhaps tell your registrars that signing of an email is the first step.

If I receive an email from your registrar that is signed then I know at least that it's from them. If I don't know that then I always have to send an email to [Matt Sterling? 01:16:00] to ask him whether this request is legit.

MARTEN VAN HORENBEECK: This is clearly something that Matt and I will have an interesting conversation about, but practically working in the Security Team at Google I actually have this problem quite a bit; that I am contacted by people that represent a particular organization and that I really have not way of validating whether it really came from them.

And then what I do is the same as what you just mentioned. Unfortunately, it's fairly high overheads so it's not optimal but the best thing then is to contact someone that you do know there to just validate and check, is this really a legitimate request? Is this person someone that I should be talking to about this particular issue?

So it's a little bit of too high overhead and there's definitely better solutions, but that is something that is one way of dealing with the problem. Thank you for flagging it.

EBERHARD LISSE: But since Matt is standing there why don't you say something about it?

MATT STERLING: Thanks. For the record I'm [Matt Sterling? 01:17:00], MarkMonitor, and I actually am the registrar of record for most Google domains. It's a fair question and we do have PGP deployed with some registry operators that we deal with. We also have a list of authorized users so that every registry knows who's authorized and who's not authorized to make changes. I'd probably defer to that.

Additionally our goal is to not make update request via email; we try to use the web portal or refer it implemented via EPP. So those would be my responses. While I have the mic, and I'm generally pretty bad about giving it up and now that you gave it to me you'll have to listen to me. I think a lot of ccTLD operators in the room have interacted with us in the past in conjunction with a number of our clients; Google and Microsoft being two of them.

And together we've made great strides in the last several years; rolling out things like registry lock, which my team has been very active in, working in conjunction with folks like [inaudible 01:18:00] in Australia and others. So I'd just like to encourage more of that cooperation and participation because as Martin said, we've seen an increase in number of DNS hijacks that are disruptive for all of us; end users, companies, registrars, registries. So we continue to be available to work collaboratively to roll more of these things out in the future. Thanks.

EBERHARD LISSE: All right. Andre had a question.

ANDRE: Hello, I'm Andre [inaudible 01:18:34] from .cz. [inaudible 01:18:35] call for the registry work but we implemented it a couple of years ago and you still don't use it. So I've just noticed that although we implemented it you have never used it. And I have a question which relates to DNS hijacking. Why are your domains not signed by DNSSEC?

MARTEN VAN HORENBEECK: The DNSSEC question is a little bit complicated and a little bit out of my area so I don't think I'm really the best person to talk to about that, but I will have to follow up with you and put you in touch with some of the people to discuss that in more detail. Regarding registry lock, if there are registrars that currently offer registry lock and you know that we are not using it then please let me know because that's something that we definitely want to look into as to why that is the case.

That might be one of the exceptions but I'll need to look into those in detail to understand why not and if that's something that would be useful to do. And that would be a conversation that the three of us would have together with MarkMonitor.

EBERHARD LISSE: Okay, one more question?

SPEAKER: [Pedre Ansad? 01:19:34], so the .eu registry, that would be another exception then? Google.eu, you're not locking that either so I should be looking at you or you?

MARTEN VAN HORENBEECK: We'll follow up. Those are things that we definitely…?

SPEAKER: And actually I have the same question as Andre; why don't you sign your domain names?

MARTEN VAN HORENBEECK: I definitely do want to answer that but I would rather get someone involved who actually has a little bit more background so that we can give you a good answer.

SPEAKER: If I'm allowed one more because basically I came for something else? You sort of mixed and matched a bit of things targeted at registries but also at registrars. I think that I saw a few of the hacks of Google.something in your list that was actually a hack on the registrar and not the registry. So you might be barking up the wrong tree for at least some of these examples.

MARTEN VAN HORENBEECK: Thank you. Just so you know, I'm really not trying to single out the registries here because this is indeed a problem of registry/registrar and also potentially people are not using some of the services that are available that actually help us protect them. In the majority of cases that I've actually looked at in the last year and half they were mainly issues at the registry level.

There were a couple in the slide deck at the beginning that are definitely registrar issues as well. But in the end the problems that we see at both the registrar and the registry level are really similar and they're not actually that much different from the security issues that the average website is dealing with.

But it's just something that we see; that there is a lot of focus today on people probing registries and registrars to probably about the same degree, looking for these types of vulnerabilities. So what we really want to do is reach out to you and figure out how we could work together to make sure that those issues don't actually affect you. And yeah, we definitely have similar conversations with registrars as well.

EBERHARD LISSE: And now the final question?

MERIKE KAEO: It's more of a comment. I want to thank you Martin for actually doing this presentation. Merike Kaeo, Internet Identity. I am very much involved and have been for over a decade in operational security issues. And registries and registrars have a really big issue in the they don't always follow the fundamental best current practices.

What Martin has pointed out are really not very difficult things to do but you have to have a process and you have to have a policy to make sure that we're all helping each other to make sure that the domain name system cannot be abused as much as it is today. So thank you for that.

EBERHARD LISSE: All right. Thank you both very much. [applause] Next is Theo Kramer for the host presentation. This one is slightly different than usual because of the way .za is set up. .za has got 18 domain names whereas .co.za has got almost 1 million. So the actual traffic and interest in how they're setting up their systems and what they're doing comes from the .co.za.

THEO KRAMER: Thank you Eberhard and thank you very much for the opportunity of the talk. What I'm going to discuss is a little bit about our transition; where we were, what we've done and where we are right now and basically a little bit about our future. A little bit about the background. Our Legacy Registration System development process and transition to EPP. Our RAR On-Boarding process, our legacy transfer initiative, some stats and what's going to be happening in the not-too-distant future.

Previously we were known as Uni-Forum South Africa and we know call ourselves the ZACR. We're a non-profit organization tasked with the administration of various second-level domains within the .za domain space, including the application and launch of a couple of new gTLDs. With around 870,000 registrations on ca.za the majority of South African domain names can be found here.

These names represent around 95% of names in .za and we believe that we are the largest registry in this part of the world and in Africa as well. Some challenges ahead that we're faced with: having transitioned to the

South African Central Registry we'll be looking at a couple of new second-levels to administer. These will include .net.za, .org.za and .web.za, which will probably become live within the not-too-distant future. And then of course our application for a couple of gTLDs. We would notify that the .africa gTLD has passed initial evaluation, as have .capetown, .joburg and .durban.

Here is a little bit about our Legacy System. It's effectively an email-based system where we accept email applications for domains using a straightforward text form. Simple operations; (N)ew, (U)pdate and (D)elete. You need to be able to configure name servers and if you can do that kind of thing you can register a domain.

One of the main reasons why we needed to transition was because this system was a post-payment system, on which you could effectively register your domain name and it would come up and you would effectively have a couple of months to pay for this. Easy to use, low technical requirements and unfortunately also easy to abuse. And we found a great deal of that for the last couple of years. And of course you also have no formal registrar relationship.

Maybe these were reasons for the success of is but also reasons for abuse. Performance statistics: you could register with a domain within minutes and up to five days if you didn't get your name servers right. 24 hours to six days for an update and 24 hours for a delete. The development in transition to our new EPP system and effectively not only to a new EPP system but to a formal registry/registrar relationship occurred over the years 2010 to 2011.

Of course the flavor of the day is EPP and we conform fully to the IETF standards. We created a system that could handle multiple concurrent policies. We needed that because the new registrar relationship required up-front payment but of course we were still running the old system, which was an email-based system, which was post-payment. And of course we were anticipating a couple of extra second levels within the South African space and these may or may not have separate policies.

The system was designed. The initial design was to separate the existing Legacy System effectively into a registrar system and separating that from an EPP system. We went operational on that during March 2011. We let that run for a while, monitoring the two systems. Effectively not difference to the registrars and registrants out there. But we monitored the system for a couple of months to ensure that there were no problems and that went pretty successfully.

Then in July of 2011 what we did was we switched our WHOIS over to the EPP database and we made the zone authoritative from the EPP system during August 2011. That's totally transparent with effectively no downtime for our registration process. The EPP system became authoritative during September 2011 when we had our first external registrar come on board. And that all went pretty successfully.

During 2012 we started with our registrar on-boarding process. Of course we had the legacy registrar or the legacy registry, which is now a registrar and we started on-boarding all the registrars. What we needed to do for that was put an accreditations system together. It is no longer a simple case of sending an email but it is now a somewhat more complex process using the EPP interface and we wanted to ensure that the registrars understood the systems, understood our policies and were able to interface through this.

We set up an OT&E server for this, allowing registrars or prospective registrars to test their systems against their systems. And once they passed a formal technical accreditation process we connected them to our EPP system, our live EPP system. And that's what that diagram illustrates. During 2013 and in fact we identified this already during the latter half of 2012 – we realized that it's actually quite a challenge to get the registrars on board, so what we needed to do was make some policy changes.

We felt that making some policy changes would provide the necessary motivation for our existing registrar base to move over to the EPP system. And what we did was in March 2013 we introduced a price increase on the Legacy System, which is a big motivator. We also introduced the notion of delegation on payment. So effectively we would publish the domain, the zone, only after having received payment for the domain.

And instead of having an expiry period of up to six months we changed that expiry down to two to three months. And what we also did was introduce a registrar accreditation fee. We had a whole bunch of registrars that had applied but they weren't doing anything so we needed to tell the existing registrars to move on and to tell registrars that hadn't made the move onto the EPP system and hadn't applied, to make a move and to get onto the system as well.

What I'm not going to talk about or try and illustrate are some of the statistics that we experienced over these periods. What we see over here are effectively the domain creates on the Legacy System. You can see this is on a month-to-month basis. You can see the average being around 5,000 domains a month and you can see around April 2013, January 2013, April 2013, July 2013 a drop on that and this is really an indication of domain registrations happening on the new EPP system as opposed to on the Legacy System.

What we see over here is the growth in domain registrations on the EPP system. Again, this is on a monthly basis and you can see we've had significant growth on the EPP system for domain registrations; growing from the beginning phases right up to this year. Domain transfers: this is an indication of our policy changes having the desired effect.

The Legacy changes happened in January or in March 2013 and you can see the registrars taking all their existing domains on the Legacy System and transferring them to the new system; an indication of the price increase on the Legacy System as opposed to a significant reduction in price on the EPP system.

And here we have some of the domains or a pie chart of the domains registered. You can see that this stat was taken on the 11th of July of this year and 45% of domains are still on the Legacy System or effectively with the legacy registrar also on the EPP system but with the legacy registrar. And 55% of domains are now with new registrars, interfacing directly via the EPP system. And the legacy registrations are shrinking on a daily basis.

This is an indication of our EPP registrars, 54% of whom are accredited and active. We now have 289 registrars. We have 117 registrars in testing and in application registrars who haven't started testing yet amounts to 126. So we've had a significant uptake on the EPP registrar side.

This is an indication of where our registrars are. Most of the registrars are of course in this part of the world but we have registrars across Africa, in Australia, in India, over in Europe, in the United Kingdom, in Canada and of course in the US as well; and this is mainly for co.za. Okay.

Here you can see our market share, our local versus international registrars. Again, stat was taken on the 11th of July. The number of domains held by the legacy registrar is 394,000 and local registrars is 433,000 and international registrars we have over 50,000 domains held by registrars that are offshore.

The growth in domain registrations: we're not seeing any significant drop. This is over the past three months and you can see our growth in domain registrations. We anticipate that by this time next year we should be on about 1 million domain names in the co.za space.

Now, let's have a little look at our challenges for next year. Of course we have our co.za legacy registrar, we have our EPP interface, we'll be looking at a couple of new second-levels; .org.za and .net.za. The new

RAR accreditation will continue, transfer from the legacy and the default RARs to new RARs will continue. We will have multiple concurrent policies for these registries and we hope to provide a seamless user experience across all second-level domains.

Our accredited registrars will of course be coming onto the new second-levels that we will be administering as well. Then in terms of the gTLDs we've got some challenges over here. We have a huge bunch of accredited registrars within the .za space but these are not necessarily ICANN-accredited registrars. And of course we'd like these registrars to participate in the new gTLDs.

So what we've done is we've created what we will call a aproxy or what we… We are in the process of creating a proxy ICANN registrar and our existing registrars in the .za space will be operating on the new gTLDs via the proxy registrar. And of course all the ICANN registrars will be operating on our systems as well for the new gTLDs. And that's it. Thank you very much ladies and gentlemen. [applause] Any questions?

EBERHARD LISSE: [Stephen Daug? 01:38:13]?

STEPHEN DAUG: [Stephen ? 01:38:25] domain registry. We live in a parallel universe. I too am doing the legacy to the EPP registry conversion and my biggest headache is synchronizing the two and producing a single zone file at the end of the day, and I'm wondering how you guys do that.

THEO KRAMER: What we did is we separated our Legacy System, we took out the registry component and we just left the registrar component. You must understand that our legacy registry was effectively a registrar as well and the ISPs or the registrars as we call them out there were just resellers.

So what we did was we created a separation and when we switched over from the Legacy System to the EPP system all that happened was that the Legacy System was sending its registry messages to the EPP system. And this happened transparently and we did that in a phase approach, we had zone publishing happening from the Legacy System first of all but then we switched zone publishing over to the EPP system. And this was after a couple of months of monitoring to make sure that there were no hiccups.

EBERHARD LISSE: Can you talk a little bit about your hardware?

THEO KRAMER: Yes. In fact I have a nice diagram somewhere on my system which will give you some indication but effectively what we've done is we're running a couple of servers. At our mains site we run both a mains system and a fail over system back to back and then we have a fail over system as well at an offsite location with a large hosting environment. And we are continually… In fact they're both a replication. All systems are a replication of each other.

Some of the things that we'll be looking at as well – we'll be moving our WHOIS system to our existing site or to our fail over site. We'll be running read operations off the various servers, across various servers and… Yeah. I actually don't have the diagram here on the screen but I can show you that, Eberhard, if you would like to see that.

EBERHARD LISSE: Any other questions? Maybe a remote question? All right, thank you very much Theo, as usual. So next will be Christian Hesselman and following will be Nigel Roberts and then Dan Timpson and then Rick Lamb. Just so that we have the order. [pause]

CRISTIAN HESSELMAN: I'm going to present an experimental service tool that we have developed at SIDN. It's called the DNS Workbench so for those who don't know us, SIDN's the registry of .nl and we currently have over 5 million domain names under management through 1,600 registrars, or a bit more than that. And we're very proud to be the largest DNSSEC in the world with over 1.5 million signed domain names so far.

And this work is coming out of SIDN Labs. It's being developed by my colleague Jelte Jansen, who couldn't be here today. And SIDN Labs is SIDN's RNDT. So what is this about? On various occasions we heard people making statements like the one you're seeing up on the slides, saying that we're looking for a specific zone file with specific resource records and specific options in terms of signing, for example, and made available for a particular name server implementation.

And then nobody really knew where they could find that or the information or that the zone file would still be available. So we figured there would be a need for a one-stop shop for a name server testing; an environment where there is multiple name server implementations available and that that environment is being well managed and well documented.

Then if you're looking for these specific zone files with these specific configurations you can actually go to one place to get it to test your name server software, for example. So that's the reason why we came up with the DNS Workbench. It doesn't really look like this by the way. And this is what it does look like. This is the overview. So the DNS Workbench consists at this point of various name server software.

So at this point it's Knot, BIND 10, PowerDNS and NSD, which are being configured from a single database containing a number of zone files. I'll get into that later. It also consists of an auto updater, that's the thingy over here, which takes care of auto updating the software on the different name server machines.

And we have documentation, which is available through our website – I'll give you the link later on and the documentation explains what version we are running on the different name server machines. What do the zone files look like? That sort of thing. And using this infrastructure you can use your own name server or perhaps your resolver or whatever else to test against the name server infrastructure of the DNS Workbench.

So we currently have this. The audio updater currently works with PowerDNS and NSD, but we're also working on updating it for BIND 10. So we're also working on including automatic updates for BIND 10 and Knot. So we think that the added value of this system is that it's a one-stop shop for an easy-to-use service tool, if you will, for a name server testing. And it supports many RR types so I'll show that on the next slide.

We have a well-documented set of zones that were in the database you just saw, and they are consistently available across different name server implementations. So as a tester you know what to get from each of these name servers. And we think that there are two specific target groups that might benefit from this infrastructure.

These are DNS developers – so they can use the responses they receive from the different name servers to test the interoperability of this software and they can also use it to discover and report on bugs that they find in name server software.

Also we think that there is an advantage for DNS operators because they can use the Workbench as a reference points for production servers, for instance, to compare the responses they see in the operational environment against the responses they're getting from our DNS Workbench.

So there are two zone files at this point. There is a signed one and an unsigned one and this is a snippet of the signed zone files. So the point to be made here is that there is a lot of research record types in there, almost all of them, so Jelte put in everything that he could find in the different RFCs. They include standard ones but also obsolete ones and also experimental ones. So this is all documented on the website.

Our current set up contains two zones, as I said; an unsigned one and a signed one. The zones are transferable with TSIG, so you can run your own secondary server using one of the servers in the Workbench as your primary and we're currently supporting four implementations, which are NSD, BIND, Knot and PowerDNS.

Some examples of how to use it – you can do it through dig for example; there is an example up there, in which case you query the Knot name server for an minfo resource type with DNSSEC enabled. Or you can use the Knot name server for instance at the primary for the secondary you're working on.

One of the things that's challenging in this particular environment is that there is potential explosion of complexity because we're dealing with different types of name servers, all sorts of signing options such as NSEC and NSEC 3, that sort of thing. We're dealing with different RR types and perhaps with implementation-specific options, so there is a lot of variables in this.

So what we decided to do was start out small and let things grow, so we're starting out with the configuration I just talked about and planning to add additional servers to the Workbench, in particular these three, and also add more zones including things like different signers and different delegation corner cases and that sort of thing.

Okay, so as I said at the beginning, this is an experimental service or tool if you want, so we're really interested in your feedback. What else do you think might be useful for this Workbench? And were also very interested to hear if it helps you in your job; if you're a developer or an operator. So please let us know when it helped and what it did for you.

Our current score at this point is that we managed to fix a few bugs regarding the handling of RR types and also a feature about the TSIG issue. So that was basically it on my part. As I said, this whole system was developed by my colleague Jelte Jansen, so if you have any technical questions I can try and take them now but if you really want to get into the gory details you should probably send him an email. Thanks.

EBERHARD LISSE: Thank you very much. Any questions? I for one will wait and see when I've got some sleepless night and start to make some sense out of it and play with it. It sounds like the more tools you get the more you must play with them and the more difficult it becomes.

CRISTIAN HESSELMAN: Well, as I said, there is an explosion of a complexity at some point.

EBERHARD LISSE: No, but don't get me wrong. I think this is a very good idea but I find when you look at these things they need 25 hours.

CRISTIAN HESSELMAN: Actually, we thought of the whole thing in order to make life a little bit easier. At least we hope that makes life easier.

EBERHARD LISSE: All right, thank you very much. [applause] So, bitsquatting. You sent us a PowerPoint again.

NIGEL ROBERTS: Thank you Eberhard. I'm going to ask for your indulgence slightly. Over the last couple of days I've come down with a terrible cough so if I break off halfway through…

EBERHARD LISSE: Excuses, excuses…

NIGEL ROBERTS: It'll be just as normal. My name's Nigel Roberts. Many of you know me. For those of you who don't I'm chief dog's body and do just about everything the .ci domain registry, which supports ccTLDs of .gg and .je – Guernsey and Jersey – and also does back-end services for .as domain registry.

Those are my contact details if anyone wants to get in touch. I've chosen a variant of a number of different TLDs. I thought we'd support our friends in the UK because Roy Arends from .uk gave me a little bit of assistance over a coffee yesterday, which I'm pleased to say [inaudible 01:5:44] paid for.

Do I just press the space button down here? Okay. Before we go any further there is a mixed bunch of people here. How many people had ever heard of bitsquatting before you read the Agenda for today's meeting? Okay.

EBERHARD LISSE: But I'd only heard of it because I read your email.

NIGEL ROBERTS: All right. A significant number but not universal by any means. Today's presentation is just a bit of in introduction. It's a bit to say there is this thing out there; it maybe be of interest. It certainly provided some interesting academic exercise for myself and a couple of other people recently. It may or may not be something that could be used in the wild, that's yet to be determined.

Certainly there is some evidence that this phenomenon exists. But what is it? What's bitsquatting? Well, I can't better the definition that you see in front of you on the screen. Bitsquatting is a subset or a weird version of typosquatting. In bitsquatting a squatted identifier is different from the intended victim by just one bit.

Now, I'm being a bit precise here. I don't say domain name because if you look into some of the other work that's been done into this it's not just domain names this can apply to. But we'll get to that. Now, first of all, why do we care? Well, it turns out that in some circumstances Internet that's addressed for a victim can be intercepted by an attacker that uses an address identifier that's different from one bit from the victim's principle address.

So if you are that principle address or you have a responsibility for is, you're going to be of interest to this. It was first described a couple of years ago by Artem Dinaburg of Raytheon. There is a link here to that initial paper. I've created a couple of short URLs for the three presentations I'll be referring to. In the [inaudible 01:56:13] bitsquatting DNS hijacking without exploitation.

But what's the theory behind it? Well, the working hypothesis is that this is caused by hardware, not humans. It's caused by devices somewhere in the Internet in the path between the person sitting typing something into his browser or sending an email or making some other query and the destination. It's not a human being typing a wrong letter.

Normally you just wouldn't care – you simply would not care. It happened once in whatever the statistic is and you just simply wouldn't care. But again it turns out that if you are a very high traffic website like Google, Microsoft, Facebook, somebody like that, this can be statistically significant.

What causes it? Well, again there is speculation about this. Heat? Overheated, overclocked servers? Poor quality memory? One suggested mitigation to this is to use ECC and higher quality memory. Background radiation or low level nuclear explosions, whatever they are – I don't know what a low level nuclear explosion is; I always thought they were pretty impressive. Cosmic rays? And of course my personal favorite: the aliens playing Halo II with live ammunition.

EBERHARD LISSE: They play Halo IV now.

NIGEL ROBERTS: Now, is there any evidence for bitsquatting? Well, yes, it turns out there is. This is Duane Wessels of Verisign. In his paper there is evidence of bitsquatting and COM/NET queries. Again, there's a URL there: gg.gg/bs-wessels. And they did some very serious work in looking at queries through resolvers.

But this is what interests me more. What are the implications? How can you use it or misuse it? How can you be vulnerable to it or defend against it? The last couple of weeks since this was brought to our attention by a colleague of Artem's, we've been doing some research on it.

And it turns out that we've just been duplicating a lot of stuff that Jaeson Schulz of Cisco has been doing. And I commend this paper too: "Examination of the Bitsquatting Attack Surface". Again, there is a link there, you can read that. It's all on the ccNSO website in any event. But basically what interests me is this question: is it being used in the wild by hostile forces and if so how do we detect it?

And here is the contention that I want to put before you today. Domain registries and registrars have unique tools and data and access that's available to them to investigate and detect whether bitsquatting is actively being pursued by bad guys. So I'd like to investigate that an find out whether that is a yes or a no.

And there is a corollary to this as well. There are other good guys out there as well, perhaps with slightly different tools. We all know that Google has access to a lot of data. We certainly know the US government has access to a lot of data.

How would you do this? Well, it's fairly easy to make a list of all bitsquatting collisions of all registered domains. It's a small matter of programming and unless your name is Roy Ahrens, who can't answer this question because I know he already knows the answer, has anybody got a brief idea of how you do that?

It turns out there is a nice little operation called "exclusive or" which helps out here. But we actually did this on a small subset of our own zone and it turns out there is a lot of them. Most of them are legitimate or ordinary typosquatting – I wouldn't say ordinary typosquatting is legitimate but it's not what we were looking for. So we need to reduce the results further with a finer grain filter.

As [inaudible 02:00:49] as possible to tools that could do this you could correlate with DNS data, registry data, WHOIS. Obviously if Microsoft has registered a variant of Microsoft and its got the same registrant and it's pointing to the same DNS you don't want to be concerned about it because guess what? It's probably owned by Microsoft.

One other possibility is IP registry data. Or you could actually look and see if the content on a particular web page is the same as on the principle address. And here's a question that's worth asking. I don't have an answer for this, I'm posing this: are there any other characteristics of domain names and registrants that could give us extra information?

So finally this is the conclusion: we have more work to do on this one. It's interesting, it's potentially an attack vector. There are other potential possibilities that arise [coughing] and a couple of us, including Roy and myself are talking about setting up a little mailing list or something like that to explore this further.

And I'd be very interested if anyone in the audience who would like to work with us on this would make themselves known, either by talking to me or by dropping me an email. And that's it. Any questions?

EBERHARD LISSE: Again I'm going to abuse the prerogative of the Chair…

NIGEL ROBERTS: I've been abused by you enough Eberhard.

EBERHARD LISSE: …for two things. Jaeson Sschulz, the person from Cisco, has agreed to speak about this in Buenos Aires on our next Tech Day. Roy was very helpful in facilitating this. And we may now see this as no problem but the Internet is growing and growing and the huge amount of data gives a large number of possibilities for failure. So I understand large scale domain names actually see traffic going haywire just because of flukes, memory things or maybe aliens playing Halo IV.

I think this is not something that we have to be too concerned about because it is relatively easy to mitigate, but it is first of all interesting and secondly it is something that we should look at just so we don't get caught with our pants around our ankles. Okay. Any questions? No? Anything from the remote side? No? Thank you very much. Our next presenter is Dan Timpson.

DAN TIMPSON: Good afternoon, my name is Dan Timpson and I am grateful to be able to speak here today so thank you for asking me, Eberhard and ccNSO. The topic that I'll be covering with you is certificate authorities and the new paradigm. And in essence what I mean by the "new paradigm" are the things that we are doing to proactively, in the industry, improve the state of SSL overall.

I'm going to make the same comment that the gentleman from Google made earlier. I'm not pointing any fingers here, I'm using this as a chance to look back at what's happened in the past. Some of you may be familiar with DigiNotar, which was a certificate authority. It was compromised and there were several hundred certificates that were misissued. The harm was potentially great. Many OCSP checks from Iran, hacking claims by the Iranian Hacker were never verified.

The response from the community was that the DigiNotar roots were removed from the browsers and the CA eventually went out of business. So the way I'd like to frame the discussion today in three main areas is first the state of SSL is stronger than ever and we've learned from these security issues from the past and we continue to incrementally improve the landscape.

We're doing this in a few different ways. Some of the ways in which we're improving the situation then: we've got industry groups such as the CA Browser Forum and we've taken measures to improve the baseline requirements for issuing publicly trusted certificates. There has also been work done with networking guidelines. There has been training and outreach programs for improved customer understanding.

And the CAs are working together to proactively respond to emerging threats that are facing us today. I'll also look at a few technologies that are emerging. There are some good IETF proposals on the table for certificate transparency and Certificate Authority authorization and public key pinning. So we'll look at those.

All right. So raising the bar in the industry. A lot of this work is self-regulated. We've got the Mozilla and Microsoft root programs, which are constantly evolving and improving the way that browsers deal with SSL certificates. The CA browser forum itself has raised the bar in a number of different ways. The ED guidelines have been revamped in 2012 and this group has also improved the baseline requirements.

They have been updated in 2013. I mentioned also that there have been new guidelines set forth for network and security controls and from a training perspective the CA Security Council has been set up and there is a wealth of information on the site of CAs and industry groups on things like OCSP stapling, the basics of SSL and it's an advocacy group that's dedicated to just expanding the awareness of SSL in general.

And of course we've got the Online Trust Alliance and some of the auditing bodies such as Webtrust and ETSI that are making many good improvements. As far as looking at things that are emerging, a good example of threats that are emerging with gTLDs, that's a particularly interesting topic for CAs as we consider the impact of ICANN releasing those new gTLDs and the threats that come along with those.

So putting it into perspective: over the past 15 years we've issued roughly 2 million certs a year. Out of those certs, the bad certs that have been issued, there's roughly around 100 per year. Most of those were from a single incident with DigiNotar during that 11-year period. Most of the breaches did not cause considerable harm and were remediated quickly.

When you look at the accuracy ratio issued each year it's very high and when you compare it to something like the Passport Office or the Department of Motor Vehicles they are not anywhere near that accurate. So has there been significant harm? It's hard to say in the DigiNotar case but the key point is that SSL continue to improve and it is getting better all the time.

Briefly on some of these networking requirements that the CA/B form has established: things like general protection of networks, supporting systems, zoning, air gapping, etc. The creation of trusted roles for systems accounts that only those that should do the certificate authorizations should have access to do so. And then there's guidelines for vulnerability and patch management, which include penetration testing and vigilant logging, monitoring and alerting.

Okay. So forward looking this is an initiative that's been championed by Google called Certificate Transparency. And what the primary goal is is to prevent mis-issued certificates without the domain owner's knowledge. And the way that this works is that CT provides published logs that can be audited by anyone. So you can go through and see if there is any foul play.

So anyone can see what CA is asserting about your particular organization. The nice thing about CT is that it's based on existing technologies that are already easily supported with industry coordination. From an internal name perspective there won't be much impact to internal CAs and certificates, as well as internal host names.

Some of the pros and cons of this technology – it does enhance the current CA infrastructure rather than replacing it. It doesn't require any actions by sites in the vast majority of cases. Some of the cons we do get hit with on the CA side is that deployment could take a fair amount of time and in order to be very useful a lot of vigilance will need to be paid to how those records get created and maintained.

Another technology that can improve SSL at large is the Certificate Authority Authorization. And in essence what this does is prior issuing a certificate a CA checks for a special DNS record, a CAA record, to ensure that that particular CA can issue a certificate on that domain's behalf. There are some drafts by Komodo that are out there for this.

Some of the key points: this is a preventative approach. Again we're trying to prevent the situation where those certificates get mis-issued. It does pull up wild cards. More than one CA can also be put into the system so that there can be multiple players there. A CAA is also a good complimentary technology to the existing eco system that's already built with CAs so there is a low barrier for employment on the CA side. And it can be phased in over time; it doesn't need to be a big bang adoption.

And most importantly it does raise the bar on CA security. So bad actors have to be able to attack the DNS and suppress those CAA checks. Some of the cons associated with CAA: DNSSEC is recommended but it's not required. Obviously it would be best with DNSSEC so that would be encouraged. It's an opt-in model and there's a perception that the CAs are trying to lock in customers.

Okay. The third and last emerging technology that is interesting to us is the public key pinning. And in essence what this is is imagine you've already got trusted roots in the browsers; there is around 65 or so trusted roots and what public key pinning allows you to do is pin or save one or two of those roots so you reduce your attack servers from those 65 down to just the one or two that you specifically want to let in.

The way it works is it can be preloaded into the browser or it can be extended into the HTTP request header so that the web server says that it will only accept a particular certificate root. And the interaction with users is similar to some of the protections that are built into the browser today, so if there's a mismatch the user can be alerted to see that there's something going on. And also these are drafts currently as well.

In the case of DigiNotar certificate pinning would have caught the problem. Once the CA was compromised, pinned implementations would have caught that right off the bat. Other pros again: it enhances the existing eco system. It doesn't suffer from the CAAs potential locking protection so there is more control from the end user side.

As far as the cons go you still have the issue of protecting against key compromise and not letting the keys get stolen. There are some challenges with creating key exchangers and there is also the impact of false positives in a hard fail mode. So those are some considerations with key pinning.

From an end game perspective where do we go from here? A lot of these proposals are still being discussed by various industry groups. There is a lot of [gutter? 02:36:40] search going on. The key point that I'd like to assert again is that incremental improvements will go forward and we'll continue to monitor emerging security threats.

Another example would be improving WHOIS. So the CAs would like to be notified of the changes that are happening from the registration side and from mitigating those kinds of threats. And then there is also the much-discussed impact of gTLD's man in the middle attacks. So we are actively monitoring the scene with those areas.

And SSL will continue to improve. Systems that are implemented and are using SSL today will gain all the benefits of these improvements. As far as nest steps and recommendations, more research in multi stakeholder collaboration is needed with the ICANN community at large, so we're obviously interested at digicert in improving the landscape and continuing to work with ICANN, SSAC and the other various industry groups.

And we're particularly interested in the CT area and have done some research in what it would take to roll CT out. There are a lot of talented people working on these issues and so the future is bright, it looks good. There are several links here that I've got for the reading and thank you for your time. [applause]

EBERHARD LISSE: Thank you very much. Question from the floor?

SPEAKER: Sure, it's [Den York? 02:18:37], [Internet Study?]. I'm curious. I agree with you on the basic premise that there needs to be some solutions to provide an extra layer of trust in SSL, and you've got a good list of solutions. But I noticed you're completely… It's not including DANE and I'm curious why not?

I ask that partly in context that we'll be having a series of presentations at the DNSSEC Workshop on Wednesday, specifically talking about this question of how DANE and DNSSEC together can add a trust layer and allow people to be able to specify which certificate they want to use in an SSL relationship, with the integrity of DNSSEC locked in on that. So I'm just curious about your thoughts around that.

DAN TIMPSON: Sure. Actually, in talking with some of my colleagues there has been some discussion about DANE and I think that DANE has a lot of merit. I think one of the advantages that we have with the CA infrastructure that's there is that we've got the processes for vetting the various organizations and individuals that apply for certificates, and so we have that eco system built up. I think that once DANE has something similar it definitely has some merit as an idea. Definitely.

SPEAKER: I guess I'd encourage you to take a look at that because what DANE provides is a way that somebody could take a certificate issued by any CA and put it into DNS for that domain and cryptographically assign that using DNSSEC. So they can specify using a fingerprint or the entire certificate but you can put a hash in there and they could say: "This is the certificate we want you to use."

So it's a tremendous compliment to the authority that CAs have already lent to that certificate because now you have the CA asserting through the processes that you have. You'd have the assertion of the certificate plus now you're bonding that with the integrity production you get through DNS where you're saying the owner of that certificate, to whom you have issued it, is now putting it in DNS.

Then it's assigning it, binding it to a global chain of trust and creating an environment where they're saying: "This is the cert issued from this CA that we want you to use." So I think you should take a look at it. I think it's got really good potential to help with… It would compliment some of the issues that are in here as well.

DAN TIMPSON: Great, absolutely. I'd like to talk more with you offline. Okay.

EBERHARD LISSE: Any other questions? Anything from remote, Kristina? Okay, thank you very much. And now comes what we've all been waiting for – at least I have –: Richard Lamb will talk about the poor man's HSM.

RICHARD LAMB: Hi, I'm Rick Lamb. I work at ICANN but this is not in my ICANN capacity. I gave a presentation that Lisse was kind enough to let me do in Prague almost a year ago I guess. You look at the DNSSEC deployment space and you see these cryptographic devices and they're either $5 and a smart card or they're $20,000 and there was nothing in between so I said, "Let's just make something." So I did.

I made it at home, I used a toaster oven to [sotter? 02:22:44] and built an HSM using what is called a TPM chip – a trusted platform module chip, which you've find in every Lenovo or Dell computer –, it's a relatively low-cost chip that does 20 signatures per second. I'll make this quick but in this talk I wanted to show what you could do when you ignore any limitations and say: "Well, okay, I went ahead and made this HSM, what else can I do with this?"

I'd like to first start with… You saw the title slide. The goal here is to eliminate key management or simplify it. My job at ICANN is to try to push DNSSEC adoption and get it deployed as far as possible. I still get plenty of people come up to me and say: "It's still too complicated Rick, it's still too complicated." I still get people; whether they're Generals or at the Pentagon. They'll say: "Here, I just want a box, just make it go away. I want a solution I can turn on and not do anything."

We're still not there yet with HSMs. Even with that we still have to roll keys, we still have to manage things. So this is hopefully a big step in that direction. So it's not so much about just the hardware. Here's my picture of a typical signing system. I really am not a graphic artist. It took me a long time to draw a picture of that little man. It really did. Hopefully that makes sense to many of you guys.

You have a signing software; DNSSEC sign zone, some source of time that you're using to say here's the start time, here's the expiration time and whatever your local time is, zone data and then some box that holds your keys – your ZSKs, KSK… I'm not going to go in to that. Your DNSSEC keys. There are many ways you could deal with this. For the root and for some large TLDs you actually pregenerate some DNSKEY RRsets. I'm assuming you guys understand DNSSEC so I apologise for those that don't I'm happy to talk with you at length at the bar.

What the signer does, what DNSSEC does, it's the regular record plus a digital signature. So the secret source is the digital signature. The way that signature is calculated is there's a hash computed on the start time and the expiration time and then whatever the data is in that RRset. That goes into the HSM, "Please sign this." You send a blob of bits into HSM and tell it to sign it. HSM comes back with, "Sure, here's the signed result."

And we've seen various presentations and various people have different implementations for their DNSSEC deployments but it's about a week; the validity time for these things tends to be a week. You need to include enough time so that if someone goes on vacation or there's a long weekend things keep running even if a signature doesn't get updated. Next slide.

Okay. See this is really bad but that's the only way I can make him look evil. That's not a bow on that guy's head, those are horns, okay? I'm sorry, I don't even know how to use a Mac. If I had a Mac I could probably draw a better picture. So now that stick figure is supposed to be the employee. One of the biggest concerns I have is the employee and that it's an internal attack. The cloud is the Internet, so you've got that.

But if you have this situation where you have either an attack coming in from the outside – because the signer is on some network; it may be behind a firewall but it's on some network – so if the attack comes from there or you have some internal attack from a bad employee, the red things are all the things that can get compromised.

Time can be set to anything. Zone data can be set to anything. Your pre-generated keys that you so carefully generated for the root and do this four times a year… Well, okay, those are just files and those are compromised. So the problem here is that the HSM doesn't know time. It doesn't care – you're asking it to sign a blob of anything.

So the attacker says: "Get me a generator signature that's valid for ten years." HSM says: "I don't care, I have no sense of time. So I sign that and I generate a digital signature that's good for ten years." You're screwed. The only recovery mechanism here now is going to be to roll your KSK. This means you have to deal with the parent. None of us want to deal with the parent.

And there have been a lot of discussions recently saying: "How often do we really have to roll that KSK?" So maybe in some designs you may not roll that KSK for many, many years, or until you see it compromised. So that's one possibility. So the crux of my talk is that this is what I'm suggesting, and if there is something wrong with my argument please let me know, because I've looked at this and I've talked to people and I think I'm onto something – but I may not be. I may completely just be smoking dope here.

So what if we put time inside the HSM? What if we put something within that boundary, that super-protected boundary that no one can get through? Because it is true – HSMs are designed well so that no one can get in it. But now we put a time source in there and we put in some other settings that say: "I'm never going to generate a digital signature longer than one week."

Now the zone data comes into the signer, the signer doesn't do a hash; it sends the raw RRsig information – here's the start of the signature, here's the end of the signature and here's the data. The HSM now has its own time source, it's own precision time source, so it compares the expiration time with that and says: "Okay, that's less than a week. Fine. I will do that for you, I'll generate a digital signature, a RRsig and I will spit it out the other end, which is great."

But what if someone then says they want to do a validity time for more than ten years, much larger than your maximum validity time? Well, now you have since you have the HSM as a time source it actually can limit that, it can either generate an error and say, "I'm not going to generate this signature," or, "I will generate this signature." So we can now always recover from this sort of compromise without necessarily rolling a KSK.

We haven't lost anything because we still have validity periods for about a week because we need that for operational reasons. So I'm suggesting that we don't lose anything by putting this sort of limitation in there and we don't have to necessarily roll the KSK. Rolling the KSK here is not going to help you because you still have these digital signatures out there that were improperly generated that have a week. But you would have that no matter what.

So I'm thinking this is a much simpler design and one in which there's going to be very little key management. This is just a further simplification. If you had a system like that I think you now can have ZSKs automatically generated, automatically rolled. If there is a compromise at some point and you wanted to accelerate the role, you could push a button, generate a new ZSK and it will go through the rolling process.

You'd never have to change the KSK. So if you did have one of those devil attacks from inside and outside, you simply could first correct the zone data, generate another ZSK and continue. So I'm thinking… I want you guys to think about this and tell me if I'm wrong. Because as you see I've actually gone off the deep end and tried to make this happen. All right, next slide please.

Okay, problem is HSMs don't work this way. They don't have time. They follow a standard called PKCS 11, which is great, we love standards. But they do not incorporate time into their calculations. I haven't found a single one. If there is somebody like that that would be great. Well, I've already built one of these things. Putting time in is not a big deal. So what the hell, build it! Right?

And there's the design. It's tamper IC integrate circuit that actually has a wire mesh around the device to detect anyone trying to get in. It's got a vibration sensor, a photo detector, it's all connected via USB, 32-bit – that's like $2.50 – arm processor, and then some low-power RAM that actually holds the keys. So if there's a tamper you kill the RAM, you shut the RAM. And this turns out to be a pretty standard way of doing things.

I got a lot of help figuring out how to build this thing from a few guys that will always be nameless, but they do these kinds of tamper evident packaging for Snowden's former employer. And they were very forthcoming in telling me how it is you go about making those things. So that was good. And then there's a TPM on the side. That does about 20 signatures per second. That's good for doing RSA calculations. I don't need that but I want a little speed.

Just to give you an idea, if I tried to do a 1024-bit RSA calculation on that 32-bit arm – and I did put code in it for there, just in case…. You know, it was cheap and I didn't want to pay for the arm – I paid for the TPM. It takes a few seconds per signature so not very good. So there it is. I have a few here. I actually built more than just two this time. So I just went ahead and… Maybe I'll pass it around. It's like show and tell, right?

Okay. So I don't know if any of you noticed in the previous picture, there is a magnetometer on this thing. The idea here is to detect tamper so it also detects movement. So what you're going to see in this device, if you hold it in a single position for a while and then turn the tamper light, which will be the third light to the right of the battery, will go on. So I'll pass it around, it's live.

There's also a little $1 vibration sensor made out of a piece of electric circuitry that generates voltage. It's kind of delicate so don't knock it around too much. Just pass it around. So I went ahead and built it. That's the back side. I actually put a smart card reader on it as well. The total parts cost for this thing by the way is about $25 including the board. So I put a smart card reader in the back.

And here are the features. It's laughable but my goal would be too… Right now there is only a small handful of HSMs out there that are [meat fips? 02:35:13] 142 level four, which is the highest level of security, one in which when you try to extract the keys it deletes everything. So there are very few of those out there. It would be really cool if I could get that. And I've spoken to people and I have some plans.

The idea here is to mass-produce these things and make it so that we can all build things that we can just set and forget and not worry about rolling keys or anything like that. It's got a TPM chip, a random number chip operator and then it's got… How do you set the parameters? If a person attacks that system; an internal employee or from the outside, what have you, if they can set the time inside this Rick HSM to ten years, then I'm screwed, right?

So that has to be protected as much as the rest of the system. That is just like the key information. And access to that, setting and configuring that is also part of this. So there is an M-of-N system in there, I wrote software for it. It's two-factor. You can do is using a smart card reader, which is what I got, using pretty cheap Acos3 cards at about $3 a piece.

Or you could use this thing that Frederico Neves here pointed out to me at a [Gallactic? 02:36:30] meeting and when I saw it I thought it was cool. How many people have seen Google Authenticator? Okay, good, a few people know about it. It's a… Did you ever see those little footballs with the numbers that change? Okay. This is Frederico's talk, but the numbers change and in order to get those things they're cheap but you've got to buy a lot of them and then you license them.

So their model is making money from you when you license the server part of that thing. Well, thanks to Google, god bless them, they wrote an App, which works on the Android and works on the iPhone, that does exactly that. So that's another way of doing that. You can get M-of-N people with this little Google App and just do this.

I modified BIND 9.9.2. It's not that hard a modification, not even for ugly BIND code. But I modified this so that it actually asked the HSM to do the full RRSIG calculation. Also the device has to have a… If anyone could just grab the hardware and load anything they wanted into it this would be really bad, right? So the device has a firmware loader that I wrote that has it's own RSA tool 2048-bit key. So it actually loads the software and doesn't execute it unless the hashes work right.

And then time. Time is the critical thing here. Time has to be accurate. This cannot be like a regular watch. Anyway, that's low cost. I don't expect you to be able to read any of this but I'm just giving you an example. If you download these things you'll see line-by-line how you go about doing authentication.

This is what Google Authenticator looks like. What you do is you generate one of these QR codes like that and all you need to do is take the camera on the iPhone and aim it at it and it sucks it in and now you've got a token. So it's kind of cool and as Frederico pointed out this is useful for so many other things. This is free two-factor authentication for the masses.

So forget about my presentation – this is something you could put in all kinds of other things as well; your registry log-ins and what have you. So it's actually a cool thing. Just to show that it worked, the red, once you ever get to that. This is DNSSEC sign zone, BIND DNSSEC sign zone. It goes to the point where it actually checks the time difference and then says, "Fine, compute the hash."

All right. And I don't know if any of you guys run DNSSEC sign zone but there's a typical output to show that it works. Keep going. This is the case where it's actually too long. So in this case I asked for a validity time that was a bit too long and in fact it fails. This is all just to prove to you guys that I actually did do this.

Okay. So some of the things I learned from talking to these spooks about making something that's a level four, something where you could put your keys in and you really completely truly trust this. Some of the ways that people look through this stuff is by using gamma rays – not just x-rays anymore –, vibrations, all kinds of strange mechanisms. So I'm thinking to myself, how do I get a gamma ray detector? I don't know. Where do we go? Ebay, right? So we start looking around Ebay for a gamma ray detector, which turn out to be very expensive devices, so… [laughs]

But it turns out zinc sulphide, glow in the dark paint. This stuff is actually used in scintillation detectors. These are the things that they have when they do these atomic… Trying to look for the results of neutrino bombardments and stuff like this. I've been in a few of these places; they just have big blocks of stuff. And I didn't know it was made out of this. It's so bright in here you're not going to be able to tell. Hold on a second.

So this is our gamma ray source but it's UV. You can see it a little bit, but if you take some zinc sulphite and shine on it it's really bright, right? So that's kind of cool. Pass this around. [laughter] All right. So even the high-end stuff, this weird stuff, in order to protect against state-sponsored actors it is possible to use the same photocell used in your TVs that receive your [locotrobes? 02:41:45]. They're very sensitive!

And I built one. You'll see it's on the board. All I need to do is coat it with a little bit of zinc sulphide. Problem of course is what sort of thing should I…? So I'm showing you mistakes. I was trying to put this in some kind of [potting? 02:42:00] compound and it didn't… This is to be done, I haven't finished this yet.

But the first thing you need to do is find out if any of this stuff is conductive. I didn't know concrete was conductive so that's what the little wires are coming out of that. They look like little bugs but, you know… Various compounds and trying to find the right high-resistance material and I finally found one. But I'd mixed a little bit of zinc sulphide with all those things so they all glow. Anyway, I found that really interesting. It's amazing what you can find on the net.

Next slide. I think I'm done? Yeah. So some of these ideas are completely new so I really do have to thank Jakob Schlyter, who did a lot of work with .sc and is also one of the key people that work with me in getting the root stuff done. Frederico Neves of course. Roy, David Miller is the guy deep inside one of the HSM manufacturers, who was just really helpful. And so many others.

But I'm hoping that what I'm leaving with you here is that when people start having… I'm hoping that you guys see that there is hope for coming up with a system for DNSSEC that is just plug and play. That we may be able to come up with something that the IT departments don't have to expend any more energy.

The problem I run into, even at the Fortune 500 companies when I talk to them, "Do you know what DNSSEC is?" "Sure I know what DNSSEC is but I don't have any more resource, I'm maxed out." And that's a barrier. We only have less than 1% of the 250-odd million domain names out there with DNSSEC deployed on them.

That's not the sole reason but it is a reason and I'd like to eliminate that reason and if there is some way to do something like this to make the problem go away that would be great. Anyway, thank you again for letting me talk. As you can tell I enjoy doing this stuff and I'm just a real geek.

EBERHARD LISSE: Thank you very much. We enjoy this stuff too. [applause] So can you do three for us please? Can you do three for .na please?

RICHARD LAMB: Oh, okay. No problem, you will get three.

EBERHARD LISSE: Any questions? Anything from the remote side? Kristina, any questions from remote?

ROY ARENDS: Rick, this is not a question, I happened to see you [working is a lot? 02:45:03] in the last few weeks. I just wanted to point out that I think the idea to put a time chip inside of an HSM is incredibly clever. That will change a whole lot. Thanks for that. It's a really good idea.

RICHARD LAMB: Thank you very much. Your beers are free tonight. [laughter]

EBERHARD LISSE: All right. Any more questions? No. Okay then. I would think we have reached the time for closing remarks and Andre Filip from .cz will do them. You can take that microphone or come up here; whatever you like.

ANDRE FILIP: Okay. I don't really have very much to say. We've had another very interesting and very successful Tech Day. It was good to get a presentation from a gynecologist and it was the first time we had a presentation from an anthropologist. I hope the next time we will also get some interesting presenters to listen to.

Again, the presentations were split between the very technical and another part was also technical but touching the local situation, which is very helpful for us to understand how this region is operating and what the problems are that this region is facing. And the last presentation was very good. That's the first time we've had a presentation about modern art here; so that's very helpful.

And I think all the organizers, among whom there is the most important one or probably the only one, which is Dr. Lisse. So thank you very much Eberhard. [applause] And thank you very much to you all because the participation today is pretty high. I'm impressed. It might be close to a record so that's perfect, especially in this country. So great.

EBERHARD LISSE: I put Rick Lamb at the end on purpose so we could keep the audience in here. All right, thank you very much. That's it.

[END OF TRANSCRIPT]