CR – Tech Day
Monday, March 12, 2012 – 11:00 to 17:00
ICANN - San Jose, Costa Rica

Eberhard Lisse: …part of the registrant to verify…

Marek-Andres Kauts: No.

Eberhard Lisse: Do the banks give an electronic readable statement information that can be harvested automatically or has somebody to transfer that information over manually into an application process to you at the registrant level?

Marek-Andres Kauts: That depends very much of the banks and in our case, we must insure that these registrars' bank fulfills these criteria. Because what is also in this FATF40+9 is that if there is a bank transaction and the beneficiary bank must verify that remitter's bank fulfills this criteria. So in our case important is to set criteria to registrar's bank and there are… For example in European Union it is very common that you can get information out of banks so bank with bank transaction confirmations automatically.

Eberhard Lisse: I mean the capturing what the remitter receiving bank give you. They can give it to you on paper; they can give you a PDF or something or an online thing. The registrars then basically have to capture this manually and put this in an application. How do they do that? Is that not a lot of work for them?

Note: The following is the output resulting from transcribing an audio file into a word/text document. Although the transcription is largely accurate, in some cases may be incomplete or inaccurate due to inaudible passages and grammatical corrections. It is posted as an aid to the original audio file, but should not be treated as an authoritative record.

Marek-Andres Kauts: No, it depends of the banks. Of course, from some banks you must take it out manually, but there are also banks nowadays who provide this information automatically. So if there is a bank transaction to your account – a registrar's account – then they send, for example, an email or use some kind of SSL channel to send you the information that… send you this bank transaction confirmation automatically. So it depends very much the bank… it depends the registrars which bank they choose to operate with and this solves the problem. And if you look to the future, of course, internet banking is developing at very high speeds, very well, I can say and this solves the problem. Did you get the answer?

Eberhard Lisse: Yeah. In my country the banks are a bit useless but some of them actually you can download your statement in an Excel spreadsheet, for example. And that you then can message into whatever information you need and automatically…

Marek-Andres Kauts: Yeah and this information is okay for us. So it very much depends the bank.

Ross Mundy: Ross Mundy, Sparta. Spent a bit of time working on DNSSEC and encountered a number of entities and organizations that were resistant to doing it for various reasons and I heard you mention in the presentation that the registrars were reluctant to accept this process. Do you have any thoughts on how to overcome this resistance, whether it's financial motivation or free stuff for them or something else to get them to accept it?

Marek-Andres Kauts: Thank you. That was a very good question – how to come over the resistance. Of course, this monetary motivation is useful, but as we are a country code domain, so we should treat all the parties equally. What we use for motivation? If we use this kind of explanation or argumentation for registrars, that if registrant isn't identified, then registrar is responsible for the domain name.

But if registrant is identified so he is in the internet, the domain name holder is in danger with his own face so the domain name holder is really responsible for the actions taken with the domain name. And so the responsibility is transferred from registrar to registrant and this was in Estonia argument to all our community and particularly registrars. Thank you.

Eberhard Lisse: We've got time for one more question if there is one. Alright, no more questions so then we can continue on. Oscar Moreno from .PR is going to speak about some not direct related DNS, or something work but security work they've been doing at the laboratory. Some of it may have application for us and I must warn you - it's a little bit deep.

Oscar Moreno: So, no I don't think it's so deep. I'm sorry - if at any time I mention something and you have any questions, please be free to tell me because my intention is to tell you about what we are doing and it's essentially security. So we are in the area of security and we have found some things that are helpful for security in both the (inaudible) and watermarking, so I intend to bring up a means of collaboration that I believe is the future, so anybody who wants to collaborate in these.

Sorry, as you can see I'm… technology can be quite down. [laughs] So how do I change to the next part? So in the one before it, it said collaboration with Dr. Dekel from Australia and myself. Dr. Dekel was one of the fathers of our area so he's the person who (inaudible) the internet as calling the term a "digital watermark."

So you know what a watermark is because a watermark is what you use, for example, in the dollar. You have $100, you go to a supermarket and give it to a lady; she will pick it up and look it in the light and what is she doing – she's looking for the watermark. And the watermark is a means of authenticating in this case that the dollar is authentic and actually a watermark in this case cannot be reproduced in the optical manner that they used for counterfeiting money in the past. So that's quite effective as a way of avoiding kind of counterfeiting money.

So we have digital watermarks and that's in the last there, as well as not digital in the right. So you can see that there is a circle and that's the watermark. Do you know that in the paintings of Chagall there is a fleur-de-lis which is a watermark? So it authenticates that that painting is a Chagall because the paper itself uses that special paper with a watermark.

So anyway, so that's what a watermark is. If it has a watermark, then it's authentic and so we have in our case we do, we have a technique for watermarks and I will tell you about what it is and so next please.

So as you can see, we have many kinds of applications in the area of the watermark. I mentioned already the authentication so you want one item authentication of content forensics. For example, it's very important in the case we have applications where you have something and it's taken to court. A case for example, some Americans, their lawyers have a lot of cases when there is a recording, for example, of something and it hasn't been tampered.

You have a video, some kind of a video and has somebody, say it was thieves or some crime was committed. And you take it to court – did somebody tamper it or not? So the question with tampering becomes a very important area where you would like to have [it authorized] so to prove that nobody tampered the… either it could be an audio, it could be a video, it could be many things.

Another example you have is the movies. Now they said the movies are immediately you put the movie out and somebody copies it and sells it all over the internet. So how can you protect that? So I was talking the plane recently and the guy told me, "Oh yeah, we have actually… I sell software for protecting the videos," he say – how do you call it – you encode the movie in such a way it's encoded, so presumably you cannot copy it.

But the people take out the encoding and then steal it and take it out. He wants to know that these copies that they make does not belong to the ones that his company is administering – that it was somebody else's. So he would like to have watermark so that when they actually play it out in the web then it will be known that it was not his.

So then as you can see, that image is very simple. In the watermark there are two simple processes. One is that you want to embed the watermark into whatever it is that you want to watermark – so that's the first step. And the second step that you need in the technology, so it's really two steps. One is you want to be able to put it wherever and be able to do it – that's embedding. The embedding is technically complicated and we have actually a [pop out] about that, about how to do the embedding in such a way that the embedding would be [sturdy].

There are ways that a hacker would like to take away the watermark and that's actually in the case you want to make the watermark very strong - that is what is called a robust watermark. And in some cases you don't care and that is the watermark is fragile. Like for example, in the case I told you before of the person that authentication, if the watermark isn't there, it's already proved that it's been tampered. So in that case you want to put a watermark which is fragile and that will prove the watermark wasn't there by somebody's tampering - so two kinds of watermark. So in that case we have a watermark that doesn't have too much power and it just makes it visible.

So actually the one before is that the correlation is very important – that's our technology is based upon the idea of correlation. The idea of correlation is that if what you are… you have the watermark and we have many different watermarks and you want to know, okay, does this have my watermark? Then you do something and test it and you find that when there is a [pick] in this relationship, you can see the pick in there so the correlation is an important component and it's explained in there.

So now the watermark – we use it in such a way where we actually put information into the… you can put all kinds of information and by putting more watermarks, then you have more information; you can put multiple watermarks. As I said, you can put a robust watermark, as well as a fragile watermark with the intention of meeting different purposes.

So we are honoring these guys with a company called Qualcomm which is a media company where they actually have an airport in Mexico City protected by their video surveillance kind of system and they're interested in every image – they want to have a unique watermark. So you would like to have a watermark – in other words you would like to have many watermarks so you would like to have a capacity that there will be a lot of watermarks.

Some of the systems, they have one watermark and that's not good enough. So you would like to have a unique watermark in every image. And these watermarks – you would like to be able to differentiate one from another. That's some of the things that the invention we have actually does.

Now I'm entering into the other area and the idea is we are using both those images, the sequences that they are called spread-spectrum sequences. So let me explain a little bit about spread-spectrum – what is the area. Spread-spectrum is another synonym for [CDMA] and spread-spectrum was a technique in science that was developed during the Second World War – 1940s – by the U.S. Military.

And what is the idea – it is very profound and very important; it's been used throughout by the military and the U.S. many, many years and it was secret until the 90s – 80s – it was a secret technology. And the idea is simple and it's still very profound. So what you do is that you have this information you have to send, this sequence of numbers, and what you do is actually you disguise in such a way that to any casual observer or eavesdropper will appear that it's noise. So you have information being transmitted which is disguised as noise, so that's really the area of spread-spectrum.

Why is that? Well, during the war, you wanted to communicate with your friends but the enemy – they are going to jam it, so that way, it's impossible to jam it because you don't know that some communication is going on. So the area of spread-spectrum is that, is the idea of spreading the spectrum so you have a signal, but that signal you will spread throughout the spectrum to make it appear as noise.

Now the CDMA which is a synonym with spread-spectrum – it means Code Division Multiple Access. Code Division, it really means spread-spectrum – I just explained, but Multiple Access is that now, the second thing you do is you like to have many users using that same space and that is the second… Why do you want that – well, you want to do that to make it more efficient. You wouldn't like to have only one user because that's inefficient.

CDMA – the opposite of that in the past –I don't know if you know – when you were trying to communicate in [wireless] and you have to ask the permission from the FCC and it was hell because I did it myself. It took one year – one full year – asking for the permission and it has to be point to point. And you use one frequency for the transmission, another for the receiving. And it was really hell.

Now CDMA – this is a technique that started the U.S. Government say, "Okay, we are freeing all these frequencies for public use. Anybody can use it." Now you are using the CDMA - many users sharing those frequencies throughout. And so that is the technique that Qualcomm made billions of dollars with that technique – the CDMA. So Qualcomm is the company that is champion in that technology. So we have a transmitter and we have a receiver so you transmit using CDMA and you receive using CDMA.

So now, what is it that we have done by the way, for both CDMA and watermarks? We found some ways to make it much more stronger – the signal – so that it would be less possible to hack it. So the problem with the CDMA is that apparently, clearly, it apparently is not… it's secret, I mean nobody can… but if the hacker is smart enough, he will listen to the communication for a long time; finally will detect that there is some real communication that it's not real random. So let me explain that a little bit because that’s why Eberhard said it’s really, really complicated, but it’s really not complicated.

What you are doing is called – it’s a mathematical area called pseudo-random sequences. Pseudo-random means it’s not really random. If it’s random, there’s no way to control it; it’s noise. If it’s pseudo-random, means that apparently appear to everybody that it is random, but really it’s not. The one who’s tending knows it’s not random; knows what are the laws and does it in such a way that it would be completely determined and he would know how to do it and there is a system behind. So that is the area – you know the system; you design it in such a way that it will serve your purposes and it will be meaningful communication going on. So that’s really what the area of CDMA is.

We have to design these sequences for these users and it’s a very exciting new area and it’s great but GPS users use CDMA too. So what we have is a method that will make stronger these signals, more secure. And this is part of what we found, it’s a method to find more secure sequences by using this method. I’ll tell you a little bit more about that.

Now this is a short code in our construction produces better sequences are the ones that they use presently in the GPS. Coded aperture imaging is another application for our technology and this is an application that has to do with x-rays and such communications and we have a patent also for that and we found a new legendre array which is very exciting and we can use it for that application.

Now let me tell you about all technology. Digimarc is a company that has the watermarks from the past. What do they use? They use again pseudo-random sequences but in their sequences, it’s the technology is such that they don’t have what is called cross-correlation. So they use a completely random… Our sequences are not really random at all, but we disguise them as random and so we can have many sequences and each of them we can know, okay, this one is different from this one.

So if we put in day 1 this sequence and this sequence in day 2, then we can know which one. So day 1 I give it to him and day 2 I give to him, then we know, okay, this is the one I gave to him. And if this one means something and this one means something else, then we're able to know, okay, this is what happened. This means this; therefore, this is what I should do.

Okay, so the usual method doesn't have good cross-correlation. You can't notice between two watermarks, so that's bad. You like to be able to make a watermark so that you will distinguish between one and another. You would like to give a watermark to him; another one to him and be able to know that this is the one that belongs to this one and this one to this one.

Now we have what is called fingerprinting and this is what I've been trying to do. What is fingerprinting? You use fingerprinting when you like to determine the identity of somebody, so fingerprinting and watermarks are completely different things because in the case of fingerprinting, it determines exactly the identity of an individual and this is really what we have – fingerprinting.

Fingerprinting and watermarks are essentially different techniques but in our case we have watermarks that are also fingerprinting. Now in the construction – and this is what maybe is kind of complicated for Eberhard – so we compose in a certain manner these sequences that will produce from one sequence to another sequence which is more stronger pictographically. So essentially what we have is a method to produce a stronger pictographic and technically what we are producing – a sequence that have much higher lineal complexities.

So this is a measure of how strong the pictographic is and so we have a method that's the heart of what we have found that this method can produce higher pictographic sequences and they are come from the area of finite fields. As I said, we use auto and cross-correlation and as you can see, the peak means that you found, the sequence that you are looking for means that you didn't find it.

The cross-correlation is the principle behind and you can see the sequences are producing it to the dimensional rates of ones and minus ones and ones and minus ones is up or down. It's kind of noise. The idea, by the way, behind, is that you want to have these principle mathematics – this is called the Law of Large Numbers. Then if you flip a coin long enough, it should be equally up or down, head or tails – both are equally probable.

So legendre array which I said before. So these are some of the arrays – how they are produced and it's highly technical, as Eberhard said. So go ahead. These are the… Actually, we have the only watermarks which are useable for medials so they are multi-dimensional. So it's the first watermark that you can do it for many dimensions. So this is actually some of what we are doing.

And so consequently, they are the first watermark that you can put in a video and lead throughout. There are all watermarks in the video that you put there in the beginning or at the end; ours are throughout the movie.

As you can see, you can use them for multi-media. Actually we can coordinate so that the audio and the video are coordinated. There is a famous case in the U.S. that there was a robbery and the guy fell down, not at the same time that the shot sounded. So it there was always the question – was that tampered by somebody or something because it's strange - it should be at the same time. No, he falls down and then the shot is heard. So if you coordinate so that the audio and the video, the watermark will be at the same time together, then that will be prevented and you will know this has been tampered or not.

So this is why now this is a project of a hospital Americans are recording. We are developing these for a hospital in Australia with my colleague from Australia and there is an emergency call the hospital is receiving and we put the watermark (inaudible). And that's a very simple application that, as you can see, it's all there and it would protect… they have this case for 10 years but the question is – is it tampered or not. And the hospitals are really worried about such things.

So the video surveillance I was telling you, a company, [Iomecian], they are interested in one watermark for every image. And you can see in there the video surveillance from [Iomecian]. As I said, they have airports for Mexico City and Kuala Lumpur.

That's again the last one. They come as community. This is project we have in Puerto Rica for .pr so this is applicable to the ccTLDs where you want to… you have, for example, a community of ecommerce and you are selling, say, paintings and you want to have a certificate saying that the paintings are authentic. So that's a very important thing when you do ecommerce. If this is authentic, it's better certified and you can put a watermark and a certificate to provide authenticity of the certificate. And that's it; so sorry if I took too long.

Eberhard Lisse: Actually you managed to 30 seconds over time which is quite excellent. I must say, being a doctor myself, I have an interest in this x-ray that Deter mentioned so we'll talk about it later. Is there any questions? I will allow one question. No, no, no, I go by the age.

Male: I don't know which way in age you mean, but okay. This looks like wonderful work. I like the… it sounds like you've come up with a set of functions that have better auto-correlation and cross-correlation properties. Is the result of this work published, *i.e.*, is it public? Can I look for the algorithms; can I…?

Male: We have three patents, as well as we have publications and it's always – for me, I'm new in this matter of patents – is always a mess – what can I publish or not. But after I clear up with the guys I know, I will be happy to give you whatever we have, but the patents are there – you can immediately look at them and I can tell you I can send it to you if you give me your email. But anyway, please, I feel honored to be working with you guys and anybody that is interested in collaborating and working together and doing something about these, I'd love to do it. It could be for money or it could be for other ways because we have both things which are for free that we'd like to share and others that we can develop a joint project.

By the way, the (inaudible), the Caribbean, are very interested in developing some kind of an (inaudible) pilot in the Caribbean using the new CDMA technique that we have. Our (inaudible) are wonderful and they're superior to what the previous CDMAs are.

Eberhard Lisse: Okay thank you. I noticed there were two more questions – Nigel and Warren – but I think you can easily take it offline so we can carry on. Mario Guerra is from the local NIC and he's going to talk to us about modifications they made to FRED – what are they using FRED about SSS EPP.

Mario Guerra: It's really about how have we used FRED – I told Marek about we are using it as a middle work tool. FRED is an EPP implementation. EPP is a tool for DNS registrations. I see it's quite popular now. It was developed by our Czech counterparts. (Inaudible) is the developer chief of that tool. Basically it is based on EPP, but it has some useful extensions, especially that one called NSSet where normal EPP use hosts and the use hosts as one of these elements. FRED was developed with NSSets which is an aggrupation of hosts, pretty useful for people registering several domains with the same DNS servers.

It has been developed in C++ and Python; used in PostgreSQL. PostgreSQL has the sequel implementation. It has several companies like modules for Apache which is a web server tool and uses SSL, Corba and implements WHOIS protocol.

Something that was quite useful for us is that it uses common-line interfaces for administering the DNS – one for administration; one for basically registration for contacts, DNS servers and domains. It was useful for us because our [telco] allocated a goal; we developed a totally homemade for our ER domains. It is under www.nic-cr webpage.

It was totally Java-based using MySQL as servers, so when we went to implement the EPP, we have to turn this develop again the DNS implementation from scratch or using already made EPP implementation. We liked the way FRED has been developed. We considered it very good; for example, it has a very good consistency in database. It also don't allow to register a domain with only name server (inaudible). That made the EPP implementation very solid which ends with a bind-like DNS tables which are text-based.

FRED has several parts for registrar zones, for billing. Especially important for us is FRED-client which is an excellent common-line interface which binds quite easily contents in the sets that is our [provision] of hosts and domains. It should register and manage this data independently. By the way, we are registrars and registrants at this point but we are planning to extend this tool to people like GoDaddy or all of the registrants around the world. We hope to use Fred-client as an interface so it is easier for the people registering domains to use this tool than to have a homemade tool using EPP.

It has common-line help. It's quite easy to transfer data to registrar in the future, so if you don't like such GoDaddy, you go to any other registrant, for example. At this point, as said, it's only one registrar now which is ourselves, NIC-CR, and we plan to use other registrars in the future. There is a good possibility that we make other people to use Fred-client as an interface and not EPP. Okay.

Now about our application – we developed about 2001 or something like that, a Java-based application totally made of free software. At this point we use OpenJDK as Java application, Tomcat, Apache and Velocity which are the Java-based tools we're using now - using MySQL5 database – totally homemade at first. We at this point use two-level and three-level domains. That is for example, UCR, ACCR for the University of Costa Rica or UCRCR for the same institution. The University of Costa Rica uses both. Tables - meaning DNS Tables were previously generated by hand and that is quite prone to error, especially if we have at this point 14,000 domains, so you can omit a period or many kinds of errors. FRED has a tool, gen_zone client, which can generate automatic DNS Tables so it is now a procedure.

We use FRED as a middle-ware, not as… when used as a webpage around FRED, but rather we use our homemade applications and we use Fred-client for generating tables. That way, we didn't have to redevelop up the application. Well, later I'm telling you that we are rethinking our application, but that's… I have that in a moment.

We have plans for substituting the [cable's] Java-based application with possibly a Drupal-based one which is a pretty interesting implementation with CMS. CMS are very… at this point are very flexible tools and are considered that it is a good alternative for developing applications because several CMSs are pretty flexible; they are very modular and it seems that we are going to use Drupal for the new application. At this point, we are in early stages of the application, but the change is quite spectacular from the Java-based at this point pretty [stiff] application and a new one that we'll use also – FRED as a middleware. And I think that's it.

Some conclusions – FRED is a software implementation. In Czech Republic it has much more than a million domains registered using this too, if I remember well. Yet, it is very friendly for end users from our point of view. What I've told you – excellent for registrars and it can be used as a top complete tool, so to speak or as a middleware as we have done. Okay, that's it. Any questions?

Eberhard Lisse: We've got four minutes for questions. Warren, you had a question for Oscar. If there's nothing else, why don't we go ahead with your question from earlier?

Warren: Thank you very much.

Eberhard Lisse: There is a question from the remote, so let’s take that one first and then Warren.

Female: I have a question. Is it possible to sign a .cr domain right now?

Mario Guerra: We just signed the .cr about one week, ten days ago. As a matter of fact, we have one client at this point which is signed also – Banco Nacionale Costa Rica Bank National Bank in our country. So that .cr is signed – its DNS is signed. Anymore questions?

Warren: So, much of the digital watermarking seems really interesting and really useful. But it seems like most of it is designed for sort of use in a court to be able to prove that this is correct, or at least that seems like one big use for it. Can you actually make lawyers and juries and people understand this or do they just hear there’s magic here; don’t worry, it’s all good?

Oscar Moreno: Well, I guess like all technology you have to have some kind of trust, so I think that to understand what is behind, I think you have to allow the technicians and the technical people to look at it and say, “Okay, yeah, this is okay; it works.” And so the idea is it works and it’s been tested. We have publications where people have looked at it and the patents. Also they have looked at it and it’s a proposal. They have looked at it and everybody says, “You have technology; it works; it’s new; it works very good,” and so essentially it’s okay.

But explaining to the lawyers exactly what it is or how you can use it, I think that’s the main thing finding and that’s what we are doing – finding applications where you can use it. But the keywords are “fingerprinting” – that you can look at it. We have a new way of fingerprinting that’s actually watermark also at the same time. So that’s essentially…

Eberhard Lisse: Nigel, you have a question – has it been solved? While we are setting up the other laptop.

Nigel: Again, this is actually quite simple. Would it be something that’s easy and convenient to apply this kind of watermarking to, let’s say web pages so that they are guaranteed to be the same in transit as they are delivered or maybe even DNS packets? Would there be some advantage over DNSSEC?

Oscar Moreno: I would say that the most you can use in the security, the better and I would say use (inaudible), I mean encode it and at the same time put a watermark. Do everything that you can because, for example, if you have a logo or you have a trademark or something, I think to put a watermark in your trademark I think is a safe thing to do. Actually, watermarks have been used by the eMarket for many, many years. What we have is a different technology that has also fingerprinting, so we have improved on the old technology. But the watermarks have been used for a long time.

Eberhard Lisse: Thank you.

Nigel: I have a microphone if somebody needs it.

Eberhard Lisse: Thank you. We're just dealing with a small technical hiccup. Fortunately Macintosh works.

Antonio Godinho: My name is Antonio Godinho. I represent the .MZ registry that's from Mozambique. We are in the process of switching our registry from a manual process to CoCCA which is a software kind of like FRED that was presented before.

So this basically is some information about .MZ registry, some history, when it was created. It is currently run by CIUEM which is the IT center of the university, so the ccTLD is still run by the university the way it was originally created. The university is a state university so it was also the first internet service provider in the country at the time; that's why it held its role of ccTLD.

I heard the Estonian earlier saying that their registry was small with 66,000 domains. Our registry only has about 3,000 domains, so this is small; it's not that big. That's why we were still running it manually up til now. We only register second-level domains so we don't register directly under .MZ but we register under .co.mz or .org.mz; others – gov.mz.net and so on.

So I've already mentioned this – all the registrations are done manually; there is no online system. There is going to be very soon but there isn't so far. We also have some secondary service using the PCH anycast cloud which is provided free of charge.

So the aim of this presentation is basically to share with the forum about the migration of our registry to an automated one; from a manual process to an automated one. So the choice for a registry system – this is the main issue here. It is not easy to select, mainly because the university does not really have enough funds to either build its own or acquire a commercially existing application.

We were also offered in the past I think possibility of running the registry by somebody else outside the country, which we didn't really like the idea, so we also stayed and decided to just run it in-house. In 2004 we approached .br mainly because they are also Portuguese-speaking so Mozambique is a Portuguese-speaking country, so we approached .br to see what was the possibility of porting their existing system to register .mz domains.

We actually went to .br to see the system and see if we could somehow use it for our own but it turned out that the .br system was really made, developed in-house mainly for their use, so it had a lot of stuff which is built into the code and anything that needed to be changed, the code would have to be changed. So this is a bit complex for us. We really wanted something that we could easily change from our system.

We did try to make it operational but after quite a while we failed for different reasons. We had all kinds of issues to run the application. By chance in one of the ICANN meetings, I came across the CoCCA software through someone that was working for CoCCA at the time and there was also a presentation on CoCCA and from this presentation, I noticed that a lot of the stuff that we needed was within the software. And this was a while back so a lot of updates have been done since then.

We looked at the software and we felt that this software covered everything that we needed to do to register domains automatically or to automate the system and not have a manual one. It also offered the option of having registrars which was something that we wanted to do because at the moment it's just one central registry; we don't really have registrars. There are some registrars that do register their domains with us but it's all done manually. So if somebody tries to register a domain with one of the registrars outside, what they basically do is they have to do a manual registration with us because we don't have any system which is integrated.

In 2008 we started testing the application but since we didn't use CentOS or Ubuntu or any other flavors of Linux that basically supports CoCCA, we were using FreeBSD on all our servers. So we tried to make it run on FreeBSD for quite a while. After successfully running it on FreeBSD, we decided that we were going to switch to CoCCA but since we didn't have any database like an online database that we could transform and import into CoCCA, so we decided we should first make this database from all the manual registrations and this work was then assigned to a group of people that would actually create this database that we could later import into CoCCA.

So this seemed to be quite a simple process, but for some reason this database was never ending, so people are just always saying that they were adding the content, but this work was never finished. So eventually the whole process stalled at the time because of this wait. So since we didn't want to wait and we also noticed that since we stayed quite a while without moving the system to CoCCA and the number of updates that had been released, we noticed then that it would probably be better to run the system on CentOS or one of the other Linux the way it was developed to run because then it was easier to do the updating.

We didn't want to have issues whereby we could not get support easily because we were running it on a different platform. So we decided, well, since we have still not implemented it, we might as well run it on CentOS. After we decided that we should do that, we decided now at the end of last year that we are going to test it on CentOS because apparently that's the operating system where they developed CoCCA. CoCCA has, by the way, changed the name to Pomoja, but we're still referring to it as CoCCA.

It was supposed to be released this year in the beginning of the year but, well, now we are stuck with the tender process for the hardware so in Mozambique things take quite a while to go through, especially if it's government-run. So we are now waiting for the tender process to be completed for the acquisition of the new hardware to run the system. And this is basically the point where we are now.

I was going to show just a quick practical session whereby we were going to just show you how the file looks like to be imported into CoCCA and how easy it is to import the data into CoCCA. Unfortunately the laptop is not showing on the big screen.

Eberhard Lisse: Do you think you can use the browser and connect to it?

Antonio Godinho: No, but then I need to put the file there.

Eberhard Lisse: We can also speak about it a little bit in Syria. I know a little bit about it and I don't want to take away his presentation, but I have a little input into the script that was written so I have…

Antonio Godinho: Yes, Dr. Lisse was very helpful in providing with the script to convert the bind zone files to have the format required to import into the CoCCA system.

Eberhard Lisse: CoCCA has a feature import file, a comma separated value file. Just click it on, click it on and then it asks for a file and it sucks it in. So basically you have to develop a little comma separated value file and I'm into (Inaudible) so I got into that and found out that you can read the zone file and then retrieve the name server and the domain entries and name servers as objects. Even if you can't program as badly as I am or as poorly as I am, it's basically a few hours thinking about it and writing this up and a few experiments and it works I think.

It sucks the file, it generates it. The point is it wants to have the name of a registrar; it wants to have the name of an [Edmon] contract; it wants to have the name of a technical and billing contact. So you have to generate one in the CoCCA [too as a solo] registrar. You must have one registrar; you must have one contact and then you use the same names as the values and then you suck them all in under the same name and then there's two ways of doing it. Either you hire students to do it for you or if you have registrants, you tell the registrants pull your domain and (inaudible) for the transfer and clean up yourself. At the same time my advice is when you tell them they must clean up the WHOIS.

Antonio Godinho: Yes, I was interested in the first presentation with the WHOIS because we also have the same problem. There isn't really any validation then for data that is input into the forms when people register the domains. There's just no validation at all. It doesn't really matter who registered the domains; the domains can be registered by anybody. So far we don't have any restrictions. It can be foreign; it can be an individual. They don't even present an ID to register domains, so WHOIS information is probably not a good place to look for validity.

Eberhard Lisse: I must say we had the same problem. We had one registrar where the management is sniffing shoe polish too much probably because the gasoline is too expensive and in the end they were just not listening. We write them and write them and write them and they just don't listen. They have basically just put themselves in as the registrant and as the registrant and as the Edmon contact so we enforced our contract and give them 90-day notice and kick them off.

And that's the only way – in the end you must have registrars, you must write a nice registrar agreement – that's my advice – write a nice registrar agreement so that if they don't follow policy you just terminate the agreement. And I think having registrars is better because you get rid of all those individuals that don't understand what a domain is; they don't understand any of this. If that means you need to have a help desk, if you can push this in the registrars, you have less expenditure as a university for these things…

Antonio Godinho: Yes, one of the main reasons to change was also the fact that you can have registrars on the system, so this was very important for us. That's where we wanted to go.

Eberhard Lisse: What my advice is what Marek was saying – make a good contract with the registrars and force the owners on the registrars. What was the real reason? What was the actual technical problems – day-to-day problems you were encountering that make you say, "No, we can't do this on paper anymore." Can you say something about that? Can you say a little bit more about it?

Antonio Godinho: Yes, one of the problems is basically the WHOIS – there is no information. So what we got a lot of requests from other registrars outside that they need to be able to access WHOIS information to see the state of their domains or if the domain is registered, if they can register. So since there's no automated system, we can't provide that information online. So they have to request it by email which becomes a bit cumbersome.

Eberhard Lisse: Any questions from the floor? Oh, then you got off easy. Thank you very much. Let's first see whether we can get the laptop to play. You should be able to see it there. When you go on presentation mode you should able to see it. Okay, our next presenter is Isak Jacobsen. He has become part of in joke these days because if I want somebody to go away, I tell him to go and talk to Isak Jacobsen who will then stand up and say, "I am Isak Jacobsen and I run .fo."

They also run FRED and I heard about that they had a web front of sorts and though I personally don't run FRED, I know several country codes are and I think it's always good to hear what people are using it as middleware or messaging it or twisting it or adding value to it and therefore, I already pushed it in Dakar that he must come again and if he was coming whether he would be willing to present and here he is now.

Isak Jacobsen: Thank you. As Eberhard said, my name is Isak Jacobsen and I am the Chairman of the Board of the Faroese ccTLD Council. And first I'll take a bit of the history of .fo and some facts about the Faroe Islands.

The first question I always get is, "Where is it?" No one knows where the Faroe Islands is but here it is. We are in the middle of the North Atlantic Ocean, approximately the same distance between Iceland, Scotland and Norway. Some facts – there are 80 islands; that's the position of the globe and that's the size of it. Size population is only 48,000, only 48,000 people, so it's not a big country. The capital is only 19,000 people. We used to say that the biggest city in the Faroe Islands is Copenhagen because there are more Faroese people in Copenhagen than the capital.

The language is Faroese and we say it's the language of the Vikings. It's a self-governed country within the kingdom of Denmark and we are not a member of EU. And the flag is called Merki and looks like that. And we like to brag a bit. This was what the *National Geographic* said about the Faroe Islands.

The structure of .fo is that the University of Torshavn, together with the IT-association of the Faroe Islands and the Minister of Industry appointed the independent FO-council. The history of the counsel – or of .fo – it started up in the University of Torshavn in 1993 and the first .fo domain name was created in '95 and the Faroese IT-Association and the Ministry of Industry elected the first FO-Council in 2001.

In the beginning it was a Navision-based administration system; it was implemented in 2001 also and the FO-Council produced the first set of rules in 2002. New regulations were implemented in

2006 and FRED was implemented in 2010, and the present council was elected in 2011.

The FO-council is appointed by the Faroese IT-Association and manned by two representatives from the private sector and two from the public sector. It's a self-owned institution governed by regulation created by the council, approved by the Minister of Industry. There are no employees at the FO-council; it's a hobby, so all of us in the council have full-time jobs elsewhere.

We have one lawyer, one marketing specialist and two IT specialists. The chairman – I - was the former president of the Faroese IT-association. The Vice Chairman is the head of the IT at the University in Torshavn and the first member is a lawyer at the Municipal Office of the City of Torshavn and the second member is a marketing specialist.

And the .fo council is a non-profit organization, not regulated by national law, renting out domains for one year at a time with the obligation to prolong the contract for one year at a time. Today there are 3,109 domain names total. In the beginning .fo domain names were first and foremost for the Faroese companies and people. You had to prove your right to a name and not abuse a third person's right. .Fo administration should check content and usage and the administration system was insufficient and required much manual work.

Now we have loosened up a bit of the regulation. You can now apply for 1.fo, a.fo, fi.fo and 123.fo, etc. We can block names of national or public interests like city names, island names and so on. And we do not check content or usage. Only a court order will make us consider whether or not we'll take down a domain name. FRED is in action and performing good and is not a labor-intensive system.

In the near future we will have a broader ns-base in collaboration with PCH and DNSSEC in collaboration with PCH and by the way, this all will be published today in collaboration with PCH and also IPv6. And we have to update FRED because we are still running the first version and we need to review the FRED-.FO coding and maybe to liberalize the regulation a bit more. And we have thought about to open up for registrars, but we are not unanimous about that.

The FO-Administration is outsourced to a local security firm. When you apply for a .fo domain name, there are two ways. You can apply with an A-application or a B-application, but first you have to identify yourself with a copy of your passport or Social Security Number. Then you have to prove your right to the name with sufficient documentation from a national or international registrar. Then you can apply with the A-application and you can now rent a name for one year at a time.

If you can't prove your right to a name, you have to apply with the B-application and then the name will be published on nic.fo and the national newspaper for a month. If no one objects with valid documentation and applies for the name, your application will be approved. The application fee for A-application is 400 dk and B is 900 dk and annual fee is 450 dk or 60 Euros or $80 US.

This is the face of .fo for the main user. There is a Faroese version and there's an English version and this is what it looks like. We have also a video tutorial down here for those who want guidance on the way. We are running FRED server 1.10.0 and I know now 2.3.20 is available and it runs on the Fedora core 9.5. Here are some statistics and some other statistics of the data. The amount of domain names is so small that it really doesn't matter.

What is FRED-FO? It's a web front end for the FRED system built with PHP and MySQL or LAMP. Using FRED client functions for updates (PHP calling Python functions); uses SQL scripts for some select PHP and PG queries. Uses external payment solution for application Wannafind in Denmark and we are not using the FRED email functions on notifications. It runs on common hardware; no special requirements and the back end administration for web frontend is also done and secure.

Some of the features of .fo or FRED.fo – register and modify user, user database and login system; different privileges for holder; technical contacts and billing contacts; user blocking and administration; forgot password function; register domain names; renew domain names and delete domain names; administering existing domain names; is the domain name available; register, modify and delete organizations; email functions; notifications and reminders, etc.; an invitation for contact changes; email history and summaries of your engagement in .fo and more features of WHOIS web interface; multi languages; logs; statistics; general settings; blacklist and quarantine administration; prices and products; external payment system; glue records administration; payment reports; possibility for free renewal of domain names; import system – well, that was fast.

To register a user, you must be a person – not a legal person but a person-person. In all our pictures we have this timeline where you are in the process of registering and an organization, after you are created as a user, you can create the organization and you can register a domain name and it’s pre-paid. And this is the summary I was talking about further up in the presentation where you can see all your domain names and you can make them… you can see all the holder and technical, the building [counters] under each of these if you push that green button there. And this tells that it is ready for renewal.

FRED-FO – this is the way it works. We use PHP function to create user and we call Python functions and we use PHP functions and MySQL database. And this is the FRED part of the system and this is the FRED-FO part of the system. And here you can see the function we are using. And that's just a bit of the PHP call for… Yeah.

Eberhard Lisse: So, any questions? The first question that I had and he already answered it to me, so I'm putting it again – is this Open Source?

Isak Jacobsen: It is but we will not put it out yet because we think that the programming is too sloppy. We will put it out in the next generation and that's one of the first things we are going to do now.

Eberhard Lisse: Any other questions? Is this… so basically you are interacting directly with the individual registrants?

Isak Jacobsen: Yes, we are.

Eberhard Lisse: Is this basically a solution for the registrar part of your combined function?

Isak Jacobsen: It is, yeah, but we put both the registrant and the registrar in the same position.

Eberhard Lisse: I know, I know, but I'm just saying if somebody… the one thing that was a bit difficult for us – there were two things in the beginning for us about FRED. First I couldn't get it to install, and it was a few years ago when we were running an Ubuntu version that was I think higher than the work it was running to and even a full day with Yagomir as a programmer in Prague didn't help much, but the drinks were nice.

But in the end, what I always felt that FRED was missing was a package for a registrar. So that the registrar can use this and attempt to do these things and interface with this whole process. How do you guys in Estonia do that? Are you also your own registrar or do you have a registrar/registry modem?

Isak Jacobsen: We have a pure registrar/registry modem and we have 41 registrars, including mine which is international companies and ICANN accredited registrars. So we used FRED in a registry/registrar relationship in a very universal way.

Eberhard Lisse: And the local registrars – how do they… did they write their own things?

Isak Jacobsen: Yes, basically there are two ways. One is to use this EPP client which is provided by FRED and as this buy-in fee or this fee to sign a contract with us is relatively low, we have also registrars who offer this as a side service and not a business by itself, and they use this EPP client. But of course, the bigger registrars – they have written their own EPP client and so of course, these registrations are done fully in an automatic way. And also concerning this question which raised before about the data – getting out the data from the bank, this is also automatic.

Eberhard Lisse: There is no… I'm asking is I'm into development and emerging registrars. Now I mean Faroe is a very small place; there will not be many emerging registrars in Faroe. If somebody wants to start doing the job and wants to register in that job, if that job were using FRED, I'm looking at it from that perspective. But the command line client – is that what you refer to?

Isak Jacobsen: Yes, I meant the command line client, yes.

Eberhard Lisse: Alright, any other questions? Identify yourself please for the document.

Ambrose Ruyooka: Ambrose is my name; I'm from Uganda and I work for the government. I'm interested in knowing if you have in the local legislation – cause I'm seeing some issues there like not allow some (inaudible), just that the displaying of an application is for some time for comments. Do you have any local law or legislation that governs the operations? Thank you.

Isak Jacobsen: No, we don't. All the regulation is done by the FO-Council, so there's no national law; we are not under any national law at all.

Eberhard Lisse: So they make it up as they go along.

Isak Jacobsen: Yes.

Eberhard Lisse: That’s good and it seems to work.

[Ope Odusan]: Yes, my name is [Ope Odusan] from .ng - Nigeria. I saw on your website you have a couple of names that are protection leased. So my question is what criteria do you use to determine what domains you put on that protection?

Isak Jacobsen: It is the FO-Council and the regulation we have made allow us to protect names of public interest. For instance, as mentioned, island names and city names and so on. Names of public interest and also names of interests of .fo… I mean NIC.FO of course, and also fo.fo and names.fo and so on – all these names that would obscure .fo we also protect.

Eberhard Lisse: I think the important point is not so much what names are in there or who decides what names, but to have a predictable process so that people know names that do like this and then maybe if we don’t allow somebody to register a name inciting our president or our founding fathers as we call it – previous presidents. We don’t do that so the names are in there.

You can try similar things; basic things. The point is that this is well documented and the stuff you can put where you want as long as people know this is the things we don’t allow, then the details are not really important. There was a question here and then Marek.

Karim Attoumani Mohamed: Thank you. My name is Karim from Comoros ALAC. I saw in your website your language is a lot of special characters, so a natural question is do you plan to add the IDN option to allow users to be easy in their language?

Isak Jacobsen: We have not implemented IDN yet, but we are thinking of it but also there we are not unanimous, so we are working on it.

Eberhard Lisse: That's a political operation. Technically, FRED was written in the Czech Republic; they have got criticals in their language, so unless you can turn around – the guys are sitting right behind you and they will contradict me – but I think it supports IDN out of the box. Almost. But that's the guy you want to talk to; he will tell you. Okay, if there's no more questions, I will give the floor for a word to our sponsor for the lunch.

Male: Sorry, I had two questions. First is the easy one – it's my particular interest because in Estonia we have a multi-stakeholder policy development process and the question which always arises which constantly is discussed is domain name price. And you have also a relatively high price, as our price is 17 Euros, but how you explain it or argue it to your community to have a relatively high domain name price? That's the first question.

Isak Jacobsen: Yeah, the price… we put the price at that level because we would stop domainers because in our view domainers, the only thing they do is highering the price for the end user; nothing more and nothing less.

Male: Thank you. And the other question – I will try to be brief and don't hold up the lunch, but the other question is regarding geographical location of your country, as your country is very remote from the rest of the world and you are, I think, very much dependent on these undersea internet cables and data cables and of course, they are not very reliable usually. Accidents could happen with them. How do solve this threat or this… this is indeed I think cyber security, [to track down] (inaudible) and how do you solve this? Do you have some kind of replica on mainland Europe or this kind of solutions developed? Thank you.

Isak Jacobsen: Thank you for the question. Yes, as you point out, we are very dependent of the sea cables we have, but we also do have a satellite connection, but of course, it's much slower and much expensive but we have a redundancy both down to Europe and up to Iceland. So yet we have not had any problems with it but, of course, it can arise, but it is out of .fo's jurisdiction. It is the telecommunication companies in the Faroe Islands that own those cables and run them.

Eberhard Lisse: Back up… we run CoCCA; we run it on [post SQL] like all of this, we run an hourly backup. We escrow it every six hours to a foreign country with the safety restriction; that's not a big deal. Find yourself a country that has a good jurisdiction that has got capable and fast web hosts and just dump it every six hours over there. Do an hourly or quarter… how many – 60,000 names – how many – 2 or 3,000?

Isak Jacobsen: 3,109.

Eberhard Lisse: Okay, we have got 2,445. This is a small amount. The audit is the biggest stuff and we back up [Poscas] every hour and we escrow it every six hours over to our host site. And for a small place the DNS is totally separate; the DNS is any-casted [catalog] – if something won't happen. It's just that you don't lose the data to reconstruct everything if you have a catastrophic failure. That's the simplest thing to do; it's not really a big deal.

Isak Jacobsen: The question – why I asked it - it's connected with the events in 2007 when the cyber attacks were and then these international connections were closed down and it's basically there are two internets. One is domestic and one is international. And our companies and also government wants to have their sites and this name solution working both domestically and internationally and also they want to make changes to the zone file. So these are my biggest concerns in talking about this domain name registry located on a territory which is remote and very much dependent on quite unreliable connections.

Eberhard Lisse: No, the registry is located where the people are. It serves the local community. It serves the local and the international internet community but the local internet community is on the islands. You must separate the DNS name server solution from making changes to zone file. There is no such thing as doing database-oriented emergency change to a zone file that you cannot do over some other means of communication if you have the primary at some location, for example, [use ISC].

If it's really a big drama, we can call them up or send them a fax or do some other means – radio – whatever it takes – communication – to actually tell them do this and this. Even in Estonia I don't think there were many instances where emergency changes to the clients needed to make emergency changes to their domain names in the DNS.

Isak Jacobsen: This was basically a theoretical scenario and, of course, it's done in this theoretical scenario and with an emergency situation, of course, and it's… this probably it could happen is very low. I don't want to hold on lunch anymore with this question.

Eberhard Lisse: I'm trying to stall a little bit because we must first figure out where it is. So I can stop stalling because we found out. I can deliver the word from the sponsor?

Nigel Roberts: Thank you. This will be very quick cause I know you're hungry. My name is Nigel Roberts. I'm the ccTLD manager for Guernsey and Jersey - .gg and .je. For some historical reasons, I'm also a Director of a moribund organization called ccTLD Services. We used to organize ccTLD meetings just before the formation of the ccNSO.

My Co-Directors are Willy Black who used to be the representative for Nominet.uk and Elisabeth Porteneuve who was .fr. We've got some money left. Enjoy your lunch.

Eberhard Lisse: And I might make one small correct – it's not a brown bag; it's a white box.

Part II

Eberhard Lisse: So now that we have all digested and ruminated, we can start again. Afternoon session starts with the usual host presentation. Luis doesn't need any introduction; almost everybody of us know him very well and he will give us a bit of an idea what they are doing locally.

Luis Espinoza S.: Hi, good afternoon. The first thing is welcome everybody to Costa Rica. I hope you enjoy a lot of our country and enjoy the ICANN meeting this time. We want to talk about a few things we implemented, maybe push it a little bit before the meeting. We hurried a bit with these things, but with help from many people we can do many interesting things. And some of these things were announced in tomorrow's speech from the inauguration. Many of the speakers mentioned about the .crs will be the sign, the DNSSEC sign.

Then I will talk about how was the process; how we implemented that and at the end the idea is to show you that could be not so hard to do it, then the managers really don't have the DNSSEC to do it.

Well, again, my name is Luis Espinoza; I'm the CTO of NIC internet, Costa Rica and we run the .cr. Well, I want to give you some instruction on planning the DNSSEC; how long it takes; some research development we did and implementation and some results.

Thinking about DNS to provide an environment where it's a good idea to have security, stability, availability, performance, reliability. How it looks DNSSEC in this environment – like one block like Anycast could be held in the performance or it could be held in the reliability or by sample, the standard of OWASP about the best practice for a code in our service because many break-ins into the systems comes because some of these practices are not followed. Then DNSSEC is a component of this environment that should help to provide many of these elements.

A little bit about planning – you will see there's many long terms, long times because it is low to take sometimes decisions but notice that it's low to make decisions. The principle we think about is we need to automate our process from our portal system where is captured of the request from our [customer ID]; the registration domains, the change is in the website. And this website is not connected with DNS.

Then the first thing we need to do is to implement that to automate and we decide to use EPP like Mario might call it before. We implement EPP and we automate the process of the different operations of the DNS, go through DNS automatically. Then we take more than two years implementing these because many things to research. Some development in the interface between our [active port] developed in Java like Mario mentioned that before to connect to FRED client and in other ways to connect to FRED that now is our middle ware.

Then we started thinking about DNSSEC and since the beginning we were thinking about to provide a trust as possible for a small ccTLD because we can use easily a marker, a small device, a secure device and it's cheap. We think about we can use some of like that before buy HSM – expensive for us – HSM. Then we want to look for affordable solution that can help us to implement the sign. We can sign without hardware; we can do it just by software but because this is about trust, then a trust you can show is better for the customer mainly. Then we want to use hardware.

Then after that research and some development of our own solution, with the help of many people here, we did a sign. We have our domain split - .cr is split in many domains by example, for commercial is .co, .cr; for government - .go; .gr; we use two letters. We have for finance sector at fi.cr. And very common we have most of the national banks in these sectors.

Then the first thing we want to do is provide to this sector DNSSEC to improve the security, especially for banks because they runs money of the people. Could be a huge interest in this topic. And that's it. What really take one of the main banks in Costa Rica – Banco Nacional de Costa Rica – the sector outside is bnonline.fi.cr in the process of signing.

Well, what do we take in account to this implementation? We want to focus on financial institutions. We want to do that because it's a good example and it's a good thing to produce awareness in the people because it's money. Focus on financial and that soupcon is small; it's only a few domains on there.

Then we want to create trust because all these systems is based on trust. Then because we want to create trust, we need to do the things in the good way, following good procedures, we look for use how to work for key generation and for (inaudible); have a policy of (inaudible) statement. And the other thing is very important for me to implement DNSSEC is if you implement DNSSEC only one part of the change, until the change is not complete it's not very useful. It's good to have it there but if you're not completely changed for the end user it's not so useable. These three main things are what's very important for the DNSSEC implementation.

Some details – we look for an important bank in Costa Rica under fi.cr to present the pilot project. We talk with the security people from Banco Nacional de Costa Rica. The first things we found the banks… from the perspectives of the banks DNSSEC doesn't result all the security; it's just a little complement of security. And they have many other issues to be aware about.

It was a good thing that they looked with good eyes this propose and we take time from them to work in this implementation. We really don't. They compromise with the product and in a few weeks we can work with them closely in this implementation. Then we really can sign .fi.cr and put the DNS records in .cr and then sign .cr. Then we need to send our DS records to IANA for the inclusion in the root servers.

Here we receive a little help because it's something like fast track to be ready for the meeting. Not fast track but they help us a little bit to everything go smooth. By now we have it - .crs are really on the root servers. The other very interesting thing about this is we are using a hardware-based solution. It's a new low cost solution based on DPM which is the (inaudible) device in most of the computers, in servers and in this area we work very close with mainly Richard Lamb who is here from ICANN. He helped us a lot with these things and it was so funny. It was a fun time trying to push this to really work because what I think is in the past is nobody do it before, using this kind of technology.

In the end we found some good solutions and very easy to implement and – well, not so easy to implement – that could help many other ccTLDs. It's mainly the small ones that are growing. And the other important element in all of these kind of technologies is the procedures. The procedures are very important because these procedures create the trust enough. If you follow this procedures you can be audit and anybody can check what you are doing is what you're saying you are doing.

And then we work too with DNSSEC policy statement, then we have a Spanish version of the DNSSEC policy statement but it's not approved right now by our Board because we don't have enough time for the Board. These kind of things take a little more time than the Anycast things. But it's the process and it's part of the process. Then we need to adjust this DNSSEC policy statement to reflect how we do the key management of our signing process.

Well, our goals – for this meeting the goal of providing DNSSEC awareness in the community – in our community, Costa Rica – we have the Banco Nacional. They embrace the technology; they embrace DNSSEC. We talked the first time with them; we explained a little bit what we were talking about. They like the idea then suddenly they start to working on their own systems to sign their zones after we have some work meetings to provide enough knowledge about what this was and how to do it.

And suddenly they called me; they were ready – 8:30 in the night – they called me. They are ready to send me the DS and this was amazing, just because they were ready for us; they were ready before we can have the .cr and the root servers. The good thing about this is they embraced the technology. This will be a lot of help to [expand] this knowledge.

And the other goal is the implementation. Now we have a signer that used this (inaudible). We have procedures that sign and resign our zones; that is integrated with all the workflow in the DNS operations and these things happen each hour for example. Yes, we have the DPS in Spanish.

How looks our system? Mario talked to you about this but I will show you a little bit more. We have a portal system that runs on the other center that runs on Java and Apache and has some credit card processing for payments and things like that. Then that system runs a process that generates all the operational changes – new domains, modification of contacts, modification of name servers, using FRED client – providing parameters to FRED client. FRED client communicates with the FRED server that we have.

Then FRED server provides to that [is hands-on], it's for zone file generation. This FRED hands-on runs zone file generation; then we have some scripts that verify the zone, the syntax, the size of the zone and many issues we have in the past, we put control things in this script. And it is constantly improved to avoid some kind of issues that we have about these manipulations of sound file by text.

Then this is without DNSSEC. After this verification we load it into the Bind and we do every load and this system would notify in the hidden server, modify to the master servers and then after that it is distributed to the secondary servers.

Well, how it looks with DNSSEC. Then we have the same scenario then we put it in the middle. The server that… after the signs have generated to file and it's been filed by this strip – it is passed through DNSSEC signer that runs with DPM technology and all this is signed. After that it's verified again and then reloaded and exported in [the zone] of the DNSSEC signer.

When hand-on, we got these files, these files one for each of our signs and after that we run another script that mainly runs DNSSEC signed zone with some parameters and generate the [dot signed]. That's easy. This script provides the ping for the DPM device.

Then this script does the [dot signed], all the dot signed files, and because these are some zones of the .tr, we need to include all the DS records of one of each of these zones to the parent, to our root. Then we have another script that runs after this first one finishes satisfactory, finishes well. We check that then this is another script does a [categorization] of all the DS files into the .cr file then runs the DNSSEC signed zone, into this from Bind; again using the open (inaudible) key libraries that is linked to the use of DPM. We produced the zone file signed. That's it. We load Bind and the rest of the process you know well how it works.

Then this is about key management. Oh yes, a simple… this is our motivation. This is the size of the scripts, just 57 lines, 31 lines. These are very small scripts that do all the process and these are [Shell] scripts. It's coordination work.

Key management – the DPM is not cryptographic processor. It's not AKCS11 by itself. It's a device that is [added] in many computers, small computers – laptops or… because it's the standard of the industry to protect some… hard drive like hardware.

\ Then some people from IBM wrote a code to use that cryptographic device to looks like a PKCS11 device. Then that so far creates all the structures necessary to gain an obstruction of the use of the device. Then the device is used like any HSM, any (inaudible) device but works different.

We have an offline laptop with DPM and we generate the keys in a safe environment, in a protected environment. At the last of the presentation we show you what is the laptop. It's not a laptop right now; it's a small PC because we don't have a spare laptop for that, but it's the same thing. Something you can buy for $300 – something like that.

Then you can see where we generate the keys in that PC then in this offline off-net DNSSEC signer we enter the signing keys then we move the public half of the signing keys to the signer… no, to the offline laptop where it's generate the key signing and we use like (Inaudible) to generate key signing key on that machine. And then sign the zone signing key with the key signing key generated in the machine offline; then move CERT and the signed DNS key back again to the [production signer], this one. This is okay with you? Anything to add? Okay, we are working together on this. Then with this system we can sign the keys.

The idea here is we keep the basic concept that we have an offline system with some protected by hardware key and that is how is work this DPM system. We found some interesting things mainly Richard did many tests then we found some things here. This DPM is a trusted platform model. This is the standard for many (inaudible). Right now almost all of our (inaudible) use this habit – you'll see it there. [It's only able to be viewed by] (inaudible) system but it's still there. It's supported by Open Source software (inaudible). It's (inaudible), not too fast. It's something like 1,034 signature by second. It's the speed of the system because it's not cryptographic the device accelerator by itself.

Then this is very important. This has a built-in hardware [RND] number because this is used for the seed for generating, so this is very important to have. If we can use it for our signer, only with the [ARIN] it is a good thing because it's there and it is a good thing to start to sign using some random hardware generator number.

Then PKCS11 interface simplified into HSM. Basically all these structures to signs using DPM are the same structures that we buy a third-party HSM because this has [attractive] layers is from Trousers provide this access.

Inside the DPM… No, this is Trousers/opencryptoki framework. Inside the DPM there is SR key – it's a special key; it's not like a (inaudible) key. It's a special key that is protected by hardware – it's only one key. It's only one key each time we set up the [wires]. We lose that PC; that key is lost. If something happens with the motherboard, that key lost. It's completely tied to the hardware; it's tied to the motherboard.

Then that key – it's only one key – is used to protect the key generated by the framework of PKCS11. All these relations – I'm not so sure about that but, mainly the concept is only one hardware key very, very well protected. It's used to wrap the other keys. Then it's secure, secure because I have a company that protects my keys and this company does hardware and there's no way to reverse that.

Well, some pros and cons about this technology – the migration, this hardware key migration, well, at the end we can do it because there's many so far out there but it's not easy. The software is provided only for Windows platform and from the maker of the chip. Some of them are Ethan Young, other one is Intel, Broadcom. Depends on the maker of the chip that provides the software but it's not easy to get from internet. The main thing is to not provide an easy way to replicate this information. It's very well protected.

Then the other thing is hardware driver support – Linux for example, has a very well drive support, DPM – as soon as I activated in the Linux then you start to see, and there's a driver for Windows too. This is low of course because it's not a cryptographic accelerator. Not all use key management framework. We need to figure out how to do these things, read all the C codes inside the framework to find out how it works. But at the end it's okay.

Some pros – it's a good thing; it's a simple procedure. You can find it someplace. You don't need a powerful PC; you don't need a new system; you can use a used system and it's free. It's a good thing.

Okay, after implementation of DNSSEC, when need to trap the signed zones to the secondary service and we found one problem with this. This was maybe known by many of you who are running DNSSEC. This [fiber], we need to change the maximum size of the UDP packets because it's by default it's limited to 500 or [12 kbps]. Then we need to modify it because it's bigger than that.

We found how to… it's very hard to find out how to back up the PM keys by self. But we found a way and we can mitigate from one DPM to another DPM the keys. And that's a good solution because that's the way I came back at me, my keys. At the end it's very simple, by the way. And the process to approve and publish the DPS will be a nightmare probably because some lawyers have to read it and well… DPM is low and in our system right now it runs each hour and it takes 15 minutes to sign all the process and publish the other assignments.

This is the location of our secondary service. I'm not talk about that. Just a little bit about bnonline.fi.cr. We have one face to face meeting to awareness about DNSSEC, that works perfect.

We have one telecom work session with Mario and the team from (Inaudible) International. And they are sign in by themselves. And suddenly they send us the DS because they are ready to do that.

Right now you can check with bnonline.fi.cr and you can check all the changes complete all the changes it is correct. Well, more or less, that's it. Thank you. I want to show you just a couple of pictures. This is a picture of the room where is located the offline KPM system. This is along there behind a door and there's the cameras. And there that is keep recording. And this is the backup of the keys in evidence box that we send to the safety box in the bank, one and the other one we keep it safe in the office. With these keys we can [grade] the keys at the KPM system. Thank you.

[Applause]

Eberhard Lisse: Thank you very much and you have got one minute and six seconds to spare. Okay, any questions? Comprehensively. Thank you very much. Our next speaker is Chris Davis, a well-known security analyst. I understand he's currently living in Ottawa but he is from Canada. And he is well-respected in the community. He is credited with having identified and taken down the Mariposa botnet – I think it was the largest one on us. And we had quite a little small problem. I'm getting the logistics sorted out but we are very efficient working; we don't have a budget and we manage to get in the end what we want anyway.

Among others, Luis from .cr-nic, we thank very much for helping out with the accommodation. He's going to talk a little bit about intrusion detection and mitigation. I don't think it's actually his primary expertise, but I'm quite sure you have got something to say. Don't contradict me before we're even started. But I'm quite sure whatever you're going to talk about is going to be very interesting.

Christopher Davis: Good afternoon everybody; my name is Chris Davis. We're going to do a quick introduction of myself and my co-presenter, Zach, and then go into the meat of this. It's not really about intrusion detection; this is actually more up our sort of areas of specialty which is DNS and how DNS is used by malware and how we analyze malware and how we take down botnets and how we stop the bad guys which is really, end of the day what we want to get to.

So as I said, my name's Christopher Davis. I currently do some work with Emerging Threats; I'm also a Fellow at the University of Toronto Citizen Lab which is a human rights organization within the University of Toronto. They had a really interesting paper out a couple years ago called *Ghost Net* which was related to some state-sponsored activity against the Dalai Lama's office and some other things.

Prior to that I've done work with a company called IPTrust in Atlanta; I started a company in Atlanta called Defense Intelligence which is where we found and took down the Mariposa botnet which at the time I think may have been one of the biggest ones in the world – sort of hard to call.

Before that I was the Director of Research for a company in Atlanta called Damballa which is an anti-botnet company – some of you might be familiar with it. Some of the botnets I've worked on – certainly the Conficker thing was a huge effort across the board and so my involvement with that was fairly limited, but you get the idea.

So most recently working at IPTrust, I got to work with some really, really brilliant people and some of literally the best in the world at what they do. So I asked one of them to come with me and present today and his name's Zachary Hanif and so I'm going to introduce Zach and let him introduce himself.

Zachary Hanif: Hi. So as Chris said, I'm currently employed at IPTrust. Before that I did my undergraduate at the Georgia Institute of Technology and during my time there I worked at the Research Institute GTRI. I've worked with Chris a bit on Mariposa and in my capacity as an employee at IPTrust, I've worked with Zeus and a number of various APTs.

My background mostly falls into large scale machine learning and large data operations using Hadoop, Cassandra and everything else in the no sequel movement. Under my belt I don't have anything super huge as far as take-downs go, but I've had around between 50 and 100 individual take-downs and sink-holed botnets.

Christopher Davis: Okay, so we want to talk about what it is that we're doing right now, today, and how this might relate to ICANN as a whole and may relate to all the registries that are here and I don't know how many registrars are here. Right now between emerging threats, the work that IPTrust is doing, work that we're doing independently, we're analyzing between 60 and 80,000 malware samples every single day.

Now this runs through a bunch of different malware analysis systems. Some of them only handle .exes, some of them handle only .pdfs and .docs – it depends on the system. What we get out of the end of that… at the result of the analysis is that we're able to see the network behavior of the malware. We actually really don't care too much about what the malware does when it gets on the system; that's AV companies; that's the anti-virus guys; that's not who we are.

What we care about is what happens when it's on your computer; what does your computer do; who does it call home to; where's the chromatic control – what we call C2 in the industry; how does it communicate to the C2 domain and/or IP address – that's really what we're interested in.

As a result of these 60 to 80,000 pieces of malware every day, we get tens of thousands of bad domains that we know to be bad. Now some of them are not just C2 or chromatic control; some of them are dropper sites where your computer would go to get a new binary before it connects to the C2, but we're able to categorize those and say, "Okay, this domain is for sure chromatic control. This one is actually a compromise site that's being used as a second stage dropper," as we'd call it.

Using all of this technology that we have and the data that we have, we're currently tracking over 20,000 active botnets. And when I say tracking, I mean actively tracking. We're able to tell that the CNC has changed 15 minutes ago; we're able to get some victim information as to how big is it, what's the growth. Sometimes – not sometimes, I'd say, but half the time, we're able to get a fairly detailed look at who is compromised. When the growth happens, we're able to say, "Oh, this is growing in a particular nation," or "This is growing in a particular way." And as we move on in the slides, we're going to talk a little bit about the type of people that are compromised.

I know that a lot of you here know how bad the threat is; how big it is, but we just want to revisit that before we sort of move into what we feel are some of the paths to a solution to this problem. Just let me make sure I hit all my talking points; I've got it written down in my book here. Yeah, I think so. Okay.

So we're going to talk about the problem really quickly. You all know this - I know you do - but I just really want to cover it in case some random person in the room doesn't know this.

Viruses don't exist – that's the easiest thing I can say. As somebody who is moderately respected in the industry and has done a lot of work, I can tell you right now what virus was 10 years ago no longer exists. Malware is designed to get on your computer, hide on your computer, hand over control of your computer – that's it. That's why we call it a compromise. It's not a virus; it's not gonna wipe out your hard drive; it's not gonna screw up your computer. Well, it might screw up your computer, but that's not intentional. That guy doesn't want to do that; he wants to hide on there, right?

Any of our solutions are generally ineffective against new threats. There's a couple of quotes here – one is from the Australian cert where it says, "Eight out of 10 pieces of malicious code are going to get in." He also said that when they initially see new malware that current anti-virus has about an 80% miss rate – that's not a hit rate; that's a miss rate – that's also part of the same quote.

This was an interesting one from Symantic that said, "Every second 14 adults become the victim of cyber crime." We actually just found that one the day before yesterday – it was kind of neat. So like I said, we all know that – I just wanted to cover it.

So the scope of the problem, like I said, when we're analyzing this malware and we're looking at these botnets, we're able to see many times who the victims are. And when we're saying that these are the victims – it's actually cutting off on the bottom; I don't know why – whatever. When we're talking about these victims, this is not, "Hey, over the last year we saw these types of people," this is, "Hey, three days ago we saw these people and tomorrow we're going to see these people."

It is more widespread than is published. We see this every single day and unlike some anti-virus companies, we just don't write press releases about it because it's so common we could put out a press release every day at, "Oh, hey, this bank is compromised; this oil and gas company is compromised; this airline is compromised; this hotel chain is compromised." It becomes sort of ad nauseum; there's no point in doing it. The problem is so prevalent that writing a press release doesn't help anything.

A lot of time you'll see in the press where people are talking about compromises and what I think the average person doesn't realize is that that compromise is actually most of the time malware related. So this is a list of the ones we know for sure that are malware related.

Now, Sony had about 12 or 13 breaches over the course of the last year and a half. Not all of them were malware related, but several of them were. RSA we know for sure to be malware related. Google's Aurora problem – we know that that was malware related. Nasdaq; the Dalai Lama I brought up because it was part of the Shadows in the Cloud that the University of Toronto did – we know for sure that that was malware related. The Mitsubishi Heavy Industries –weapon developer defense contractor – definitely malware related. United Nations – we've seen multiple compromises within the United Nations. Recently saw the International Olympic Committee and the list really goes on. I mean we could just go all day on that.

So we're going to talk now about current response which I think I'm going to do a little bit on and Zach's gonna jump in on as well. The current response, your first line of defense, is your anti-virus. We talked a little bit briefly about how it's not really effective and I'm going to explain why.

Normally I'd have a whiteboard or I would have put a graphic, but I kind of did these slides at last minute. So I'm gonna just try to visually explain it to you by just moving. So right here is what we call A Day – this is the Author Day – this is the day that the bad guy wrote the malware, okay? Then I'm going to move over here – so that's A Day – and this here is Zero Day. This is the first time we've seen it in the wild.

So the bad guy wrote it; he starts to compromise people; this is now when we've seen it in the wild. It's hard to say how long that is. It can be as little as 48 hours; it can be as much as a week; sometimes – very rarely – longer than a week. And then over here is Signature Day. This is the day that the anti-virus company wrote the signature. That on average, is a week between Zero Day and Signature Day. I'm not saying always; some AV companies are great and they're able to get a signature out right away. But the problem is that when you have 60,000 pieces of malware being released into the wild every day, you know, you can only have – as my friend, Paul [Wall] from Georgia Tech would say, "You can only have so many children chained to desks writing signatures." You just can't keep up and there's no good automated way to do it which is why you're going to see all your anti-virus companies now coming back with generic trojan. It doesn't have a name; we know it's bad, but we don't know what the hell it is.

Now what is really interesting is that the bad guy – he knows this; he has all the anti-virus products in his little lab; he's checking his malware against them to make sure that it can't be detected by the various AVs, which actually they have a term in the underground in the malware community and the bad guy community where they call it FUD – which we call Fear, Uncertainty and Doubt; they mean Fully Undetectable. And they sell FUD trojans for large amounts of money.

So here's your Author Day; here's your Zero Day when you see it in the wild and over here – Signature Day – which let's call it awake. Now the bad guy knows that that's gonna take a week so three days after he sees the first sort of report of it in the wild, he has it call home and update itself. Now he has to write an entirely new signature, right? So what we end up with is what we call a Constant Zero Day Window where the signature is being put out but the anti-virus companies are always a few days behind the bad guy updating his binaries.

Essentially AV becomes very effective at removing old malware and it's completely ineffective at removing new malware. I'm actually gonna let Zach talk on a little more beyond this point.

Zachary Hanif: So as Chris mentioned, anti-virus solutions have serious problems dealing with new and ever-changing threats. We've had some reports and we've seen it ourselves that over the past year – for the year of 2011 – we've definitely seen anti-viruses be able to detect probably somewhere between 60 and 80% of all of the historically published malware. The problem with that is that that's a test that has been done over an entire year's worth of malware after all of those signatures have been pushed out and updated into the anti-virus systems.

What this effectively means is that anti-viruses prove to be a fairly effective immunization tool to prevent old infections from getting in, but they have a serious problem keeping new binaries – new kinds of infections – from getting in in the first place.

Also up here we've got a few kinds of responses to this sort of problem, one of them being IDSs and IPSs, most famously done through Snort, TippingPoint – all sorts of network security mechanisms along those lines. And they are effective to a certain degree. The issue with them is that the primary thrust of the community behind an IDS or an IPS is designed to prevent network intrusion. They have very few tools, very few signatures that allow them to detect egressing compromises.

This could be things like data exfiltration, reception of command and control, messages from a CMC server – things along those lines. And the reason for this is not because they are ineffective tools, but mostly due to the fact that every rule you implement on an IDS or IPS system leads to additional weight on that system. It takes a little more CPU, it takes a little bit more memory, slows the entire process down. Many network administrators will not turn on egress rules and they'll simply focus on the ingress rules that are applicable for their networks. So while they can detect internal malware infections, we don't usually see large numbers of network administrators paying as much attention as they probably could.

Finally, we've seen some success through court-ordered take-downs. We've seen some success through NXD and other mailing lists. The problem with these is not a problem of effectiveness or reliability; it's just primarily a problem of scale. We could have every judge in the world writing take-down requests every day and we still wouldn't be able to keep up with the flood and the fact that we'd be issuing court orders that frequently is not a world, I think, anyone wants to live in. Likewise, NXD is a highly effective tool but there are only so many people who are on it; there are only so many hours in the day for which people to effectively initiate a take-down.

At the true end of the day, the problem even goes beyond the actual effectiveness of any one solution. These solutions are only helpful is someone is sitting behind an IDS sensor on the exterior of a network; if they have up-to-date and current anti-virus systems. At the end of the day not all users have such things, certainly not everyone has an IDS system. There are very large numbers of users who don't have up-to-date anti-virus signatures and the problem continues to spread despite these tools.

Christopher Davis: Thanks, Zach, you touched on kind of all the points there. For those of you that aren't familiar with the NXD mailing list, it's run by Andre Ludwig who is now with New Star, I believe. Great guy. It's a small, closed mailing list but it is a great place to sort of get researchers and registrars and/or registries together.

So we want to talk about how to fix the problem and we have this infrastructure where we can analyze all of these pieces of malware every day, where we can pull out thousands of known bad domains and we can do this with a very high level of efficacy and accuracy. What we're lacking is the ability to move that information into the right people's hands to get this stuff shut down.

And so what we're proposing today is that there should be created a public benefit non-profit domain clearinghouse for malicious domains. I know this idea has been bounced around, at least in some form for a long time within the various communities, including ICANN. I know that Jeff Moss has mentioned it; I know that Rick Wesson with Support Intelligence has mentioned it; my friend, David Dagon at Georgia Tech has talked about it for like two and a half years. Paul Vixie – oh my mic is cutting out; sorry.

So we're just gonna go through what we think are the steps to get this done and we're not gonna get up here and propose that we're the people to do this; we'd certainly love to be involved in it; we can certainly contribute to it, so as we go through and explain this, when we say "we," we mean the royal "we," not us personally, alright?

So as I said it has to be 100% public benefit; it has to be non-profit; it also has to be set up in such a way that people cannot profit from it because there is a market for victim information; there's a market for things like large amounts of malicious domain feeds and there's nothing wrong with that market being there but that's not gonna fix the problem. It's just going to make some people some money and a few people are going to get parts of the information.

So emerging threats – Matt Jonkman in coordination with myself and some other folks in the industry – are certainly willing to back this idea with as much malware as we can give it – the malware analysis students we can give it. We certainly need ICANN's backing in this. Community support – I just threw up a few people I was thinking about. I think ISC would be a good player in this; I think David Dagon out of Georgia Tech for those of you who know him would be amazing. Rick Wesson of Support Intelligence would be great; Alice's Registry would be awesome for this. There's many others that I could go through.

The primary goals here are we have all this malware; we can analyze it. The key problems start to come in when this has to be essentially if we're sending 1,100 take-down requests to a given registrar on a given day, we can't be wrong – not once. And I know that there's an entire political problem there that I'm not even touching on - on getting 1,100 domains over to a registrar but…

So identify, analyze, validate, confirm – this has to be sort of the key tenets to be able to pull this off properly. The idea then would be we have all this malware; we know what domains they're talking to, then what do we do with them – we gotta move them somewhere. Now we could do that *via* maybe using EPP; we could do it *via* maybe just changing NS records. There's a good paper published actually by ICANN which is I think on the next slide that talks about some of these problems.

And then once we have these domains hopefully sinkholed where we can now start to enumerate and identify the victims, we need to work on notification to try to clean it up. And that has to be done in coordination with a large number of bodies – CERTS; certain non-government departments; certain government departments at certain times. And so the idea here is that we would be able to notify them and provide remediation information.

If you look at Conficker, there are some lessons to be learned from Conficker, but it was an amazing group effort across the board with ccTLDs everywhere, gTLD operators everywhere – it was amazing to me. And then once we were able to sort of sinkhole and identify the victims, we have to remove in a coordinated fashion the malicious domains from the registry.

I just want to that I touch on my points. So the other important thing is law enforcement coordination. Law enforcement has to be involved in this. The more bad guys we get arrested, the better. Taking the domain away from the bad guy is great but he's just gonna go start another one – he's just gonna go write another piece of malware. And if we do this properly, we'd be able to get a large amount of evidence that we're able to gift-wrap in some sense and hand it over to law enforcement and allow them to initiate an effective and quick investigation and hopefully get the guy behind bars as opposed to just taking his toys away from him by shutting down his domains.

So the offerings to the community to the registrars and registries to other people we need to involve in this – be it NGOs, government departments, law enforcement – would be that we'll provide a daily bad domain feed with no errors. So this domain is related to this malware sample; this MD5 or MD5s – it communicates currently to this IP, this port, the malware type, the whole bit – it would have to be zero error.

We have to be able to take in transfers of those domains because if we have to get a court order every time we want to get a domain shut down, we've all lost; it's over. What we can pull from this information – and we already have some – I don't want to say who it is – but we have some people onboard already with the idea of creating a bad actor database where we're able to say, "Hey, this is the email address this guy used; here's the IPs he logged in from when he registered the domain." We're gonna put that into a secure database, make it available to law enforcement and also available on a per query basis to registrars and/or registries - and I'm open to ideas about this obviously – is so that if you get somebody registering a domain, they can go check against the database and say, "Hey, that email address was used for this type of botnet three days ago."

The way that we do the analysis on the malware, the way that we produce the information has to be peer reviewed; you can't just trust a single or two individuals to do this – it has to be agreed up experts across the board that are saying you're doing it the right way.

One of the key benefits here is that for registrars especially, a large percentage of bad guys, when they register, command and control domains, use the default name server – the default authoritative name server – with the registrar. So GoDaddy has a ton of traffic on its name servers of bots looking up their camatic control. I don't know what it's costing them but traffic-wise, it's probably pretty high. And we can get that off of their pipe and put it onto the clearinghouse pipe. Do you want to touch on how we handle all of that data?

Zachary Hanif: So obviously our main goal here is to clean up large sections of the internet that are heavily infected by various pieces of malware and everything else that goes along with it. To affect this we have to make sure that we have no false positives. A handful of false negatives is not necessarily an issue because we're not vetting domains as good; we're simply vetting domains as bad.

So obviously the largest effort we're going to have to do here is insuring that we have zero false positives and we plan to do this through the use of very large big-based mechanisms. The vast amount of information that's available from the malware analysis systems that are currently out there allows us to have extremely complex and extremely refined machine learning models that allow us to determine whether or not a domain is truly bad or not bad.

This is not a perfect solution as anything dealing with statistics, there is a variable amount of error. The point in it however is to insure that we don't have so much information every single day that a human would be completely overwhelmed. The idea behind this is to use a machine learning model to winnow down all of the possible bad domains on a daily basis and then at a final step use a human eye to insure that we have no mistakes.

It's a fortunate thing that things like Cassandra and Hadoop have come along, mainly due to the fact it allows us to cross over a vast amount of data very, very quickly and in a very scalable fashion. It allows things like active streaming feeds of bad domains; it allows things like very complicated feed requirements to be easily compressed down and reduced into smaller, much more manageable files that can be pushed to an individual who has control over various malicious domains.

So we've obviously started to discuss some of the technical challenges involved and the first obvious technical challenge is to determine whether or not a particular domain is a true camatic controlled domain or if it is simply a compromised domain that happened to be hijacked by a malicious individual.

Going through this kind of material is complicated and it is somewhat time-consuming at the start up but the models currently exist that allow us to very accurately determine whether or not it is a true compromise or if it is a simple hijacking of someone else's previously legitimate domain and we will be relying on that to a great deal. There is an extensive amount of published work that already exists that is public. There is also a large amount of private work.

Some of the better known public works are *Notos* and *Exposure*. These follow the general principles we're going to have to utilize to effectively deal with this problem and scale.

So the other major concern is from a technical and remediation perspective, the other major concern is dealing with how you want to identify individually affected users. Obviously malware analysis will allow you to find the command and control points, find where the owners of the botnets are coordinating their botnet width, but it is a significantly harder task to find individuals who have been affected by these pieces of malware. And probably the quickest and most effective way to do so is through use of extensive sinkholing.

And that brings us to our second technical challenge which is to find a way to craft individual sinkholes and policies behind them to insure that we have accurate detection of actually compromised users. For example, we don't want to do something like just simply saying everything that's on a particular port is a compromise. Obviously, otherwise, basic port scans would set flags up all day long and that's not acceptable.

So those are the two largest technical challenges for actually hunting down and beginning to identify the affected users and affected domains. There are two other major technical challenges behind that which, to some degree, we're going to need input on.

Firstly, we need to come up with a way to rapidly quickly and add bulk **trends**? for large numbers of domains from the current owner of a malicious domain to some manner of clearinghouse as we're proposing. Obviously we're not registrars; that's your world. You have the most experience in that particular area and we're coming to you and asking for your input on this matter.

Obviously it has to be fast, it has to be simple and it has to be very easy to correct as needed. Obviously there's every effort being made to insure there are no false positives, but sooner or later something will creep in and we have to be able to reverse the transfer as quickly as possible and this is a particular area of expertise where I personally don't have a lot of knowledge, but obviously the individuals in this room do.

Probably the final technical challenge would be dealing with victim notification remediation. The obvious concern here is victim privacy. We don't want to reveal any information that shouldn't be revealed and we certainly don't want to neglect… notify an individual for fear of privacy concerns. For this matter, Chris has got a lot of experience from his work with the Mariposa Working Group and he can go on at some length about that.

Christopher Davis: So the primary thing when we're talking about notifying victims – and I actually want to touch back on when Zach was talking about sinkholing before I go into this. This is whether you're involved or not, if anybody in this room is involved in doing something like this in the future or even just sinkholing a given domain to anybody – including law enforcement – the biggest thing that you have to do is not just set up a server that listens and allows the victims to hit it because you're going to get a massive amount of false positives. You're going to get security researchers; you're going to get people port-scanning just random garbage on the internet. If you ever like just turn on a sniffer on the internet, you see how much garbage is out there.

The biggest lesson we learned with Mariposa and Active X intelligence when we were setting up the sinkhole is that if the connection to the sinkhole does not exactly match the connection string required to communicate with camatic control, it gets dropped and not recorded. So I just want to make sure. Like I know Zach touched on that, but really important.

Okay, so victim notification privacy – when we were taking down Mariposa, we tried to sort of call people as we saw it. They were compromised, particularly in Canada cause that's where we're stationed. And we called the Canadian banks and we called some of these larger corporations and the response that we got back was 1) “Who the hell are you;” 2) “Did you hack us; are you trying to extort money out of us; I’m calling the cops.” - just really good responses.

Every now and again somebody would say, “Oh hey, we found the machines that were compromised. Thanks very much.” So when we’re talking about privacy, this needs to be done *via* a known body to the area. So whether it’s a CERT; whether it’s a government department; whether it’s an NGL; whether it’s a larger governing organization like ICANN – I don’t know what the solution is to that, but I can tell you right now – hiring three people to pick up the phone and call everybody is not going to work.

Okay, so moving on to the next slide. Special challenges – the biggest thing that we have to get out of this is we have to get the registrars and the registries to buy into this idea. If everybody says, “No, I have to have a court order,” we’ve lost and I think we all know that. If a court order is required, then you’re going to get X number of domains and they’re not going to be very big, when we’ve got 60,000 pieces of malware being released in the wild every day. We have tens of thousands of new domains that are malicious, that are stealing people’s information – stealing your grandma’s credit card number; stealing even way more important information than that.

So we need to come up with a community way to provide proven zero error information that can get this stuff taken down. And so this will not work at all without registrar/registry buy-in. Community-wise, of course, we need the support from ISPs, from CERTS like I was talking about before and also large industry partners. I know the guys at Microsoft; I know some of the guys at Google, I could go on about that and certainly draw up support for that, but there’s a lot of people we’re not touching here.

So the first steps – what we're proposing today is that we will immediately start working on providing a per-registrar and per-registry feed for free to every registry in the world and every registrar in the world; a daily list of malicious domains – the malware; the domain if we can, the IPM port that it's currently communicating with; if we can, the botnet family. So Verisign will get all the .coms; GoDaddy will get all the .coms that relate to GoDaddy. Do you follow me on this?

We're going to provide that as soon as we can. We think that we can probably get this done within the next 60 days and start getting it out. We're going to start actively supporting Snort and – for those of you not familiar with Emerging Threats, Suricata, IDS project, we're going to start providing focus rule sets on detecting compromise within environments based on the data that we have around the malware.

All new TLDs – this is another thing. If we had the ability to analyze malware like we do now 15 years ago or 20 years ago, it would have been a lot easier to keep up with what the bad guys were doing than to go back and say, "Hey, there's 100,000 domains that we know to be bad right now," or a million domains – whatever it is. So we're gonna say all new TLDs.

Beyond this registrar/registry feed, we're actually going to start actively monitoring new TLDs as they come online. So name your… .green, let's say. If that's open to public registration – which many of these will be – the minute we see a piece of malware, the first one, use that TLD, we're going to notify everybody involved as quickly as possible and then stay on top of it.

We also want to offer this to ccTLDs that may not have the resources to do this and all of this work right now – this is our volunteer work. Nobody's getting paid for this but we know we can do it; we can do it in our spare time and in our off hours and we'll offer to do that starting right now. Do you want to add to that, Zach? Okay, and I think that's it, so we're just down to questions and answers now.

Eberhard Lisse: Who wants to start?

Nigel Roberts: Thank you. Nigel Roberts from .gg and .je registry. I've got a couple of concerns about this. First of all, I got a bit of a sense here of a courtroom in a Wild West whereby the posse is saying, "If we wait for the guy to be convicted, we've lost so we may as well go out and hang him now." Now, my anti-spam, anti-malware credentials shouldn't really be in doubt but I'm not an apologist for the bad guys – far from it. But I'm very concerned for the fact that when you take somebody's domain away from them, you're taking their property away from them.

Bad guys own property; good guys own property and our society has come in such a way that we have something called the Rule of Law and it seems to me that we're trying to set up something here that would completely bypass that. And I'll just give you one final thought. Can you explain to me what a malicious firearm is?

Christopher Davis: What a malicious firearm is – that's very good. I think the difference… I don't disagree with you. I actually have a lot of the same concerns as well. the biggest issue that I see is that if law enforcement which is under-staffed and overwhelmed by the problem, has to go get a court order to take down domains that we can, through a peer review process – not just a random hang 'em high judge in a courtroom – say, "We know this to be malicious to 99.9% certainty," that's kind of what we're striving for.

You're touching on exactly where the weakness in the plan is and the problem is that I don't see a solution for that problem. I see a way for us to perhaps take care of 90% and there's 10% that's going to be an area where it has to be a court order; it has to be maybe done differently. The malicious firearm point was very good but if a domain was registered three days ago, it is being used as command and control today and three months from now, it's being used as command and control for three different pieces of malware, there's a pretty good change that domain has no legitimate use.

The issue is where we have what we call mixed domains where a domain was used for something good and is now being used for something bad and is not compromised or the domain was compromised and we can't take that property away. So we have to be accurate; we can't confuse a compromised domain or a mixed use domain with an exclusive camatic control domain. And that will take work and it will take probably a lot of discussion and conversation on the way to do that.

Eberhard Lisse: I mean, the point you're making is if it's a legitimate domain, you must avoid taking it down. If it's an illegitimate domain – and I use this word with forethought – who will come and complain? They won't because if they do, they can complain and we'll have a chat with the Royal Canadian Mounted Police while we're doing it to find out what's going on there.

But still, domain names are property. If you register it in a registrar and I take it down, I run the risk of getting into, in several federal U.S. jurisdictions in the United States, domain names have been described as property and the use of an email address under such a domain name has been charged with $5,000 plus punitive damages. So this is at least under federal jurisdiction in several districts in the U.S., you have to be very, very, very careful. I've also seen a Canadian judgment in a different context that also says its domain name.

In practical terms, if you take a purely command and control domain down, who is going to complain? The point is if somebody comes to complain, you then and because you have made a mistake, you run into serious, serious trouble. Okay, the next question comes remotely.

Female: Hi, I think there are two questions from Antoine. The first one is, "Why would you want to transfer domains? Isn't changing the .ns enough or do you want ownership because you can?"

Christopher Davis: No, we're totally open to the idea of just changing the NS. The idea is that a lot of registrars would be overwhelmed with the number of NS changes and as somebody that doesn't want a registrar – I can't say this for sure – but my experience is that both domain transfers happen all the time – bulk NS changes may be a little more complicated, which is the only reason we put that up there. But we are totally… I don't know the best way to do it – that's your space far more than mine.

Eberhard Lisse: I would think that name server changes need more verifications – is it correct; is it not? If you just change the ownership and only the owner can make these changes, then these changes are not going to have to be evaluated. But again, again, if you take down a criminal domain or if you take down a firearm, who will complain? It's just a point that you must be really, really, really, really, really accurate.

Christopher Davis: The accuracy is the most important thing.

Eberhard Lisse: Never mind that technically speaking, you're still taking away somebody's property. Now one can regulate this with policy and if the registrar's policy says if you do that, it's violating our policy, you can still wait for somebody to come to complain. But technically speaking, I'm not sure that no law, no law enforcement can take away somebody's illegal firearm from somebody. Can only be law enforcement; it cannot be you, even if we know he's going to kill somebody with it. Sorry.

The monopoly to use force against an individual in a state lies with the government authorities, the police and so on. I'm not opposing; you expose yourself to serious liability. There was another question from remote.

Female: Antoine also says, "I don't see how domain equals malware. Could you please explain. I understand command and control but I think you only see domains as the only hammer."

Christopher Davis: So in this context… so that question is correct in the sense that not all malware users' domains… for its camatic control or uses DNS to locate its camatic control, but it's about 97% of modern malware uses DNS to locate its camatic control. It's actually a lot easier to get an IP address shut down than it is to get a domain shut down.

So if the bad guy uses IP or IPs to directly communicate from the bots to command and control, that's fine and that's a totally different process to get that shut down. But if 97% use DNS and they have one or two domains or maybe more auto-generated domains… I mean there's an awful lot of different ways or routes to go there, but most malware uses DNS.

Stephen Deerhake: Hi, Stephen Deerhake, AS Domain Registry. Is it my understanding that you're proposing that the new gTLDs get onboard with this at the front end of their life spans and if so, are you proposing that ICANN include compulsory language in the agreements with those new gTLDs?

Christopher Davis: No, I'm not suggesting that the new gTLDs get onboard right away. What I'm saying is that we'll start providing them essentially a monitoring service for free, let them know as new stuff in their space gets used by the bad guys. Beyond that, I'm suggesting nothing else. That's a really interesting question though and you made me start thinking about that.

Morgan: Hi. Morgan – I work on the Google Response Team. I actually really liked your idea. There's obviously things kind of like this already, just less legally encumbered. For instance, like the Mozilla Chrome AV… I mean sort of virus safe browsing lists and stuff like that.

The problem that I have is the identification of something actually as malware and as malicious cause you said this is gonna be peer-reviewed, right? So I happen to know exactly how much time it takes to reverse malware and even if you do it at scale, and you have like a giant malware reversing farm, then you actually get into deciding exactly what constitutes bad behavior cause you have software which does malware as things, talks back to domains, updates itself – all that sort of thing. Does that mean we read all the EULAS in terms of service and software to make sure that… So how are you going to do this identification?

Christopher Davis: You're hitting on what we spent hours and hours talking about and that's great. It's just something we couldn't even start putting into a slide deck. You're exactly right. And so there's going to be stuff that you might consider to be spyware-esque or I don't know, adware – whatever it is. And we're getting samples; it looks like it's malware; smells like it's malware but in fact, it's not.

So we would have to work with the AV companies which we currently do to merging threats to get our samples in and they have false positives as well where they're saying something's malware where it's not. That's a great question and again that is… we don't have all the answers. But I would like to buy you a beer.

Female: Garth Miller in the Adobe Room wants to know, "Will the proposed system have an API?"

Christopher Davis: Actually I'm gonna let Zach answer that but…

Zachary Hanif: So at some level all systems would have to have some means of communication between obviously the registrars and the proposed non-profit. It wouldn't have a proper API so to speak as for a public use, public reporting and public confirmation as far as allowing just random web users to come in, makes requests and attempt to influence the process as a whole.

More likely than not the actual delivery of the malicious domain names, domain names that have been determined as malicious, would be over flat file dumps, probably CSB format, something along those lines.

Eberhard Lisse: The reason why what Garth is asking us is the Chair of CoCCA, the organization that has been writing CoCCA tools [Premier], one of the Open Source registry systems. FRED and CoCCA are the most eminent ones. So it would be very helpful if the registry system had an API to you guys. That's what he's asking.

Zachary Hanif: There's certainly no reason why we wouldn't be able to implement an API to fulfill any needs that a registry would have. There's absolutely no reason that we wouldn't be able to do such a thing. That's entirely within the bounds of technical possibility. It's just simply a matter of policy who has access to the ability to request this kind of data.

Eberhard Lisse: But that's a matter that can be solved not only by policy but also by procedure, by protocol…

Zachary Hanif: Of course.

Eberhard Lisse: … that you can say, “Okay, you must register, for example, .na on CoCCA; you must register the IP address.” Like we do it with EPP – you make sure that whatever private and public keys get exchanged so that you know if a request comes from this and this thing is legit, and you talk to it – if it doesn’t, it doesn’t.

Zachary Hanif: Authentication would obviously be a serious concern. It would have to be implemented in any kind of system.

Eberhard Lisse: But what I’m saying is it might be a cool idea to think about to look at the Open Source registries – FRED and CoCCA too. If the programmers want to build this in that you sort of accommodate this and that, then it is for each ccTLD manager using the system to then establish a link with you and write and an agreement so that you understand each other and that we know what we’re on. But the programming interface – that can be organized between two part… between the programmers in New Zealand and you guys, without each individual ccTLD member having to get involved in that.

Zachary Hanif: I agree.

Eberhard Lisse: That would be a cool idea.

Zachary Hanif: It would.

Mario Guerra: I wonder if that could be… Mario Guerra from nic.cr. My question is I remember when Conficker was distributed. I wonder if this could be implemented in a similar way that, like Conficker. I think it is possible not to complicate it if… but it is a wide scale, much, much more… much, much bigger than Conficker, of course.

Christopher Davis: Yeah, I brought up that I thought Conficker was a very good effort; I thought that it was amazing to me how many ccTLD operators were onboard - it was almost everybody if I remember right. I think if you were to talk to Paul Vixie or Rodney Joffe about what they thought about Conficker, I think there's lessons to be learned there and I don't disagree with them.

Rodney Joffe was involved with me in the Mariposa Working Group and he said that he thought that perhaps we may have learned some lessons from Conficker and gotten a bit better. But no, I don't disagree with that. I think that the issue needs to be how do we – when we have maybe 10,000 Confickers – how do we deal with that? And so this is at least some steps in that direction, but like I was saying before, we don't have all the answers but that's a very good point.

Warren Kumari: Warren Kumari, Google. So I really like this idea and I really hate this idea. So I like this idea because it can actually make a difference and because I know people in the community and I sort of trust you to be sane and do the right thing.

But there are two big problems with this. The first is not everybody knows you and knows the community so they have no real reason to trust you. And the other one – which I think is more worrying – is while I know and trust you now, how do I know that five years from now or 10 years from now you will still be same and won't to tend this big scary cabal of take-down folk?

Christopher Davis: That's a good point and so when we were presenting this at the beginning, I said the "we" is a royal we. I'm not suggesting that I run this; I'm not suggesting that Zach runs this. This needs to be peer-run; it needs to be certainly the reason we suggested non-profit is at least in Canada, you have to have a minimum number of Board members that can, you know, kick out the people that go crazy. So you can't just sort of have a single cult or personality running this and that's really important.

I'd love to be involved but in what fashion – I don't care. But I think that there needs to be checks and balances and you're dead on with that because anybody that does this – that's the concern.

Eberhard Lisse: Actually it's not even necessary. If I get irritated by you too much, I just turn the connection off. I'm just saying, that's… It's not that once we connect to you, we are tied into this forever. If we find out that it's giving us too much grief, we can just turn it off. So I don't think the question that Warren posed is really a relevant one because we can always disconnect.

Dmitry Kohmanyuk: Dmitry Kohmanyuk, Hostmaster, Ukranian ccTLD. I just don't see how the service like yours can be handling common used domains. Say I have a Twitter account that is sending URLs which point to kind of a stream of [bot] control centers or have a Tumblr blog which is used to post such things or have… you know, Tumblr would be like user name@tumblr.com or URLs or any kind of this – and essentially you can just… okay, can I shut down twitter.com if 10% of its users are using, their accounts use botnets? What's like a threshold, you know?

Christopher Davis: Yeah, that's an excellent point. Dynamic DNS is a really good point too where, you know, if we have – I don't know; let's just… I'm gonna point out Sam at change.ip who's a wonderful man – but if I have, I don't know, five or six 2LDs and then I have 20 or 30,000 3LDs, you can't go shut down the 2LDs. It's the same problem. We're not going to be able to solve all of these problems. What we can solve are the bad guys that register a domain – a 2LD – to look after their botnet. If they're using a Dyn DNS provider, we've gotta go to the Dyn DNS provider. If they're using Twitter – and it would depend on the way they're using URL shorteners – I mean, that's an entire conversation to have at the bar. But there's a lot of different ways to approach it. That's a very good point though. Thanks.

Eberhard Lisse: Alright, thank you very much. That was quite an interesting presentation. And… Maybe I just spoke too much and it gave up. Our next presenter is Emre Sezglner from Turkey - from .tr. It's just my unfamiliarity with the language that I can't spell it right. We continue to struggle with the technology for a second. Just be with us for a second.

Emre Sezglner: Hi everyone. This is Emre Sezglner from nic.tr, the ccTLD for Turkey; I'm the second name there; the first name is my boss, Mr. Atilla Özgit. This is a very brief presentation concentrating on the matter of ownership of a domain and the technical issues related with that one.

The .tr delegation was made so it's pre-ICANN registration. The organization here is the Technical University, Middle East Technical University in Turkey. We are proudly serving the Turkish and the whole internet community for quite a while with .tr. Here's the domain name numbers registered totally under .tr – it's a steep curve.

The numbers will be coming so I won't bother anyone now. So we had two sort of second level domains. We decided at the beginning of the delegation that we should use the ccTLD notions so we have .com.tr.; net.tr; ort.tr and everything goes under those. There's no flat avc.tr; we do not register such a domain name.

The primary issue is when we started we also decided this was a bold decision that we require documents that proves the domain name should really be registered to the applicant. So the second level domain names on the left are documents-required domain names and the other ones on the right are totally free. You can just go on the web and enter your credit card number – it will be up and running in one hour.

So next slide please. We provide validity of the owners by looking at those documents. The documents are trademark documents; Chamber of Commerce registrations and such. So we get a copy of that directly sent to us. That is done with ordinary fax or through the web. We also have kind of validation process through the web services with other organization.

For example, if you are asking for a name – surname - .com.tr and you enter your I.D. No., that's checked through the web service and you do not have to send any documents. That makes things easier, I think. We have the registration software of our own that is written in-house, but all other software around is mainly free. We're running on Linux; we have the (inaudible) flavor.

Actually the (inaudible) software is very helpful for us. As I said, the software is developed in-house. It has various legs and arms but mainly the user interface is PHP. We have the registry/registrar model running and we use electronic pulse to collect the money.

So the documents come in and they're also managed with in-house developed software – very simple document archive program. There's a software site and the [Scandisk] is converted to PDF at the moment of being sent. If it cannot be converted, it is reported back to the user there's something wrong with that. It can be GIF, PNG, JPEG – whatever they like.

So this leads to a perfect paperless office and I think we're doing very good with this one. Sometimes you have to go back and check a document is really valid or another one sends some document to us and requires the same domain and you have to go back and check what it was. So this program – this archive – leads to a manageable accessible resource.

And we tend to use the web services instead of the hard copy of electronized digitized versions of the documents, so we are working with patent offices and such, so you only need to give the registration or application number and all will be done at the (inaudible). So this is an interesting part of the process. At the beginning we decided to require documents but perhaps it will be very different if we started today. We wouldn't do that perhaps, but when you do this, it becomes the dependency and all the community embraces that and it turns into a tradition. So we have the documents-required conduct here, second level, and all other second level domains that do not require documents and here is the numbers of registration.

We have 290,000 domains at the moment totally and 200,000 of these are com.tr which require documents. So there is an old way of registration, so people like to use com.tr because it means something. It means that that company, that website, really belongs to people that says he's the one. So some sort of SSL like thrust is being built up with this one.

So if you use this sort of domain – the document-required sort of domain – you know that the domain is yours and it will be yours forever, so nobody will come up and say, "You do not have the right to use that domain," because it was [commanded] at the first place.

I'm saying that there is a dispute resolution process in .tr.cctld which is done by the DNS Working Group. It has members of ISPs, universities, lawyers and such. So it is totally independent community. They come together and decide the disputes. But the dispute number is very, very low because when you first register the domain you have to prove that the domain belongs to you. So most of the problems are resolved at the beginning.

So I'll try to answer your questions. Perhaps you can change all the way to another direction. Please go at.

Eberhard Lisse: Any questions? What database are you running it on?

Emre Sezglner: At once we were running Oracle but it easily can be adapted to any free database.

Eberhard Lisse: Now, I like the idea of it's a paperless office. For our system we are trying now also too. For example, when we send out invoices or so, we try to make a PDF, we send it out, but at the moment we archive it in the email system. That's a bit cumbersome and we are busy sort of to put it in the database so that we later can keep it all by one domain and we can look what documents we have.

If we have… people don't have to provide documentation, but sometimes when we're not sure we ask them for found… we call it the founding statement of the company and then they scan it and send it in and where do you put it then? We don't want to put it in a folder; we want to put it with the domain electronically paperless office. I think this is a very cool idea. Is the software Open Source?

Emre Sezglner: Not at the moment it's not, but the reason is mainly because it is very specific to needs of dot.tr – all the ongoing processes, and so probably would not be interesting to any other ccTLD.

Eberhard Lisse: You probably need to speak Turkish to read the source code.

Emre Sezglner: Not at all. We like to use English commands.

Eberhard Lisse: Alright, anymore questions? Or any questions? Thank you very much.

Emre Sezglner: Thank you.

Eberhard Lisse: Ed Lewis from .us is next.

Edward Lewis: My name is Ed Lewis; I work for New Star; we run .us and we operate a couple of TLDs. The title of my talk is *The Relationship of Registries and DDoS.* It's not DDoS hitting registries; it's the relationship because there are different ways to look at how DDoS will impact a registry. I want to highlight one of them in here.

I was asked by our representative on ccNSO to give a presentation theme of security beyond DNSSEC. And I'll tell you – I've worked with DNSSEC since the beginning. It's solving a smaller problem than registration; it's just following what DNS does. There's a lot more to security than DNSSEC for registry, no matter how much is said about DNSSEC.

And so we sat around and we though - what should we talk about and my first topic was taking down of malicious domains, but that's already been done. So we went on to DDoS. DDoS is another topic that seems to be making the rounds. It's a general purpose problem around the internet and it has a special impact with registries which should have a special place in the whole world of DDoS. So I'm going to focus on DDoS's in particular security issue.

I'm going to talk about DDoS, kind of define it. You can define it in a lot of words or a few words. I'm going to take a few words to talk about it. And then I'm going to talk about how a DDoS might hit a registry – what it means to hitting a registry – I'll talk a little bit about that. But then the next line is, "How can a registry be an unwitting accomplice in a DDoS?" It's something that we see evidence of and I think everyone should be aware of what's happening out there that registries can actually help DDoSs along without knowing it. And then finally, how else can the registry be involved in DDoS protection, not just for them but for their community that they serve.

So DDoS starts out with DOS – D-O-S – denial of service, meaning I can take out a service. I can either block access to it or just crash the server – that's a very simple thing that's been around for eons. DDoS is distributed version of that. That means that what I do to take out the service, the DOS is sourced from many places on the internet.

When it's sourced from many places, it becomes hard to track. If I have one network and I see the source and destination of a bunch of packets, I can say, "They're causing you problems," and just take them out. But if a pack has crossed boundaries, I have a hard time going back through ISPs and finding out, you know, "Can you get the source of this data?" Because no one actually sees the whole (inaudible) DDoS back to back or end to end rather.

So the first D in DDoS makes this much more complicated and it's a very hard problem to crack. Why can it exist? It exists for a few reasons and one reason is that on the internet we use the client-server model. Client-server is not the only way to do networking - peer-to-peer is the other one. But client-server lets us put all this responsibility on the servers do a lot of work and clients tell them what to do. And so we have a lot of clients out there who then become taken over, become part of botnets and they can flood servers with requests and tie up the servers from doing work for anybody else.

Also we use a lot of very lightweight protocols like UDP which is send and forget protocol. I can send something to somebody else and they have to receive it and react to it and they may not come back to me. I might send a false return address to somebody and have it go somewhere else. So these are two of the reasons why DDoSs can actually be launched. We can't get rid of this – this is very important. Client-server is a good thing; lightweight protocols are a good thing. We can't get rid of what enables DDoS and there are other factors too, but these two leap to mind… leapt to mind when I wrote this.

Why do DDoSs exist? What's the incentive? We know there's monitoring incentive – there are people that want to take out a service – either steal something, stop someone else from making money – it can be very oriented that way; or it could be ideological. In fact, that's what the more recent trend is – people want to prove a point like, "I just want to take out people who don't agree with me. I want to take them off the net; I want to silence them."

When it comes to fighting DDoS, the motivations are always gonna be there; we're never gonna stop motivations. We can't do anything about that but some of these motivations will help us understand how to deal with that DDoS when it's happening. That's about as far as you can go with that.

Where are registries exposed to DDoS? No. 1 thing – they could be a victim of DDoS. I could target them and I could target any one of the public services they have there – the web, the DNS, the WHOIS server or whoever WHOIS is going to be replaced by; registration interface and so on. They can also be part of somebody else's DDoS against somebody else – the reflection attacks that come out.

The other exposed surface too is that registries tend to know more about the internet than anybody else that they serve. You tend to hear more of what's going on there; you may know about things that are going on because we sit in security and operation of the internet. In fact, you want to see the internet flourish in your area – so you're gonna keep your eyes on what's going on, so you're pretty much tuned to things. And so again you can be the victim, an unwitting accomplice or you could just be a forecaster of DDoS.

So I'm going to kind of focus down now more towards DNS because DNS is probably where you're going to see most of all of this DDoS activity. Registration systems tend to be pretty well protected because you know who your registrars are. Most of you don't have wide open registration interfaces. You don't have email problems as much. You have spam and all that, but that's not been the big thing, and so on.

But DNS is generally where registries really have the biggest exposure to DDoS events because that's where the lightweight protocols are. On the other hand, registries tend to be really good at DNS. I do some scanning. I go through the internet looking at the way people run DNS and the TLDs almost uniformly are perfect at running DNS. They have the capacity, they've got attention to this, they've got well run servers. So generally a DDoS launched towards a registry doesn't seem to make a lot of difference – it just gets absorbed or just treated some way.

But I do want to make you aware of a way that registries get used in DDoS event and then finally also, just by simply knowing what's going on, you have the ability to tell other people what's coming down the road.

Just one slide. In case you are a target – I probably should just put a slide on that. Treating DDoS is pretty simple theoretically. Just get rid of all the bad packets quicker than they arrive. If I can get rid of bad things before they… if they come in to me every second, and I can get rid of them in half a second – it's not DDoS – there's not enough of them.

That's simple to do. The hard part is knowing what's a bad packet and that's the biggest issue. And in fact, I've had like 50 slides on

DDoS on here but I'm not going to go through all of that. There are a lot of techniques to treating this. We use overcapacity, but then of course you can image the botnets get bigger and bigger and more and more capacity. Anycast is cited as one of the big helps in DDoS protection. What that does is it compartmentalizes the attack and we see this in Anycast clouds we have.

When a DDoS is happening somewhere in the world, usually just some pair or some group of our servers are getting hit because that's what's covering that part of the world. Botnets are regional. The botnets may not be regional, but their impact tends… seems to be regional in what we see. We do have global botnets.

And finally if you do scrubbing and filtering of packets, there are services that do that. They'll take a traffic in and look at it to look for signature of DDoS and drop all the bad stuff. Generally this is not… most registries already do this. This is kind of what you have to do to just be seen on the internet and I don't think it's a big problem.

So the unwitting accomplice topic – I think this is a little more serious. There's a thing called reflection attacks. They've been seen for quite some time – I don't know how many years since I've heard about them for the first time. What this counts on are two things – one is an attacker will go out there and get a whole bunch of sources of data and have them throw requests to something which will then take the requests and answer back somewhere else. They reflect the attack off of these TLD servers or servers of any kind out to somewhere else.

Now that really doesn't make a whole lot of sense to me because if I can just launch data, why do reflect it? The answer is amplification which is the next part of this. I can ask simple questions of a DNS server, get back very large answers and that's the real win for the attackers in this kind of game.

The return address which you see… a packet comes from somewhere, you answer it, you send it back to that address. That return address is not the real sender – that's the victim. You're not the victim; that's the victim, even thought it looks like the attacker to you. So with distributed systems, it's just a little less detectable.

This is my graphic slide; I don't do these too much. This kind of tries to illustrate what a DDoS looks like and the colors are kind of faded but across the top I had red colored or pink colored bad guys. This is the botnet across the top. These are supposedly the real addresses – 10/8 here I used as my example network – and they all say they're coming from 192, 168 something or other. They send queries down to my DNS in the middle there and it sends all the answers down to the real 192, 168 at the bottom of the slide. That's kind of picturing the reflection of focusing all of these individual sources down to one place down there.

The other thing that happens here too is all those top black arrows might be so small you don't even notice them. They might be such a low data rate you don't even see the traffic coming from all over the place; you don't really know they're coming from all over the place. No ISP sees it, but you see all of it come in as small little queries and they go back out the bottom.

One of the examples that doesn't involve registries but kind of illustrates the point of amplification – there's an attack where the query is for all of the data that ISC.org, the domain name ISC.org. That domain name, because of DNSSEC and IPv6 and other things that they've been adding to that domain name, it's a huge amplification – 165 times. The request for that data takes about 24 bytes, the answer comes back almost 165 times bigger than what you asked for. That's a huge gain for whoever's trying to launch this attack.

Now for the registry you would think that it might be kind of obvious that if some address is saying, "Give me information," and you're answering back to it with the same information over and over again, you kind of catch on after a while. Like if a real name server out there asks you a question and the answer is gonna be this and there's a TTL of one hour and they come back in five minutes and ask again and again and again, they're kind of overstepping the bounds. They're asking too many times – that's very suspicious.

But there's a value in non-existent names here because what I can do to make a registry less suspicious is just ask for random names that you don't have. In fact, there's no registry in the world that has more names registered than unregistered. I can go through even .com now and generate random names and I'll find I don't know how many – I haven't calculated in my head how big it could be, but there's a lot of possible names in any registry and if I keep changing that string, you're not gonna see the requests as being repetitive. You're going to think it's just another person looking for something. You might see after a while they're scamming. I mean a lot of the simple attacks would do alphabetical scans – an increment by one, ask again; an increment by one, ask again. There are just many things to ask for so it's a little harder to stopping being a part of this.

The best thing to do here is to start going through logging and start looking for this kind of scanning and generally there would be two things to look for. One is one address constantly getting these queries – again, they're not the attacker; they're the victim. And then this weird pattern of access – that's kind of what has happened.

What you can do for it basically is be on the watch for it. Just the signatures for this; look out for this kind of attack. The reason why I'm harping on this is that I know that there's a lot of activity in this area right now. Reflection attacks have been big for months now. We've seen one persistent one where there is someone attacking a lot of victims so it's hard to catch up with them cause we don't know who they're going to attack next but we know who they are. We don't know where they are cause their return address is false. But we know some fingerprinting of the attack; we know it's probably the same thing. And we've seen some of the patterns when they're active to help give us some information about where they are in the world. It seem like it's a manual event which is kind of still helping us – not fully automated.

The second thing is once you detect this you can start scrubbing and you can filter out these things. You have to be aware of when they stop cause you don't want false positives; you don't want to filter out good traffic, so you have to go back and forth over time.

So basically those slides were about the accomplice part. This is the next part which is being the forecaster and to shift your focus here, you know a lot about what's going on on the internet. The fact that you're even here, you're hearing a lot about these attacks, all these possible sources of bad behavior reporting and so on around the world. There may come a time when you say to the people around you, "Hey, we found someone doing some bad stuff. You want to get ready." We've heard this before already earlier today. The intended victim might actually get suspicious of you. They think maybe you're about to launch an attack on them and this has actually happened.

Not too long ago there was a group that was planning some attacks and the attack planning wasn't too closely guarded a secret and it got into the hands of a ccTLD. And they tried to go out and warn all the people involved that this was coming. Now even though the ccTLD was tied into the community, there were still elements of the organizations involved of being intended victims who didn't know who the ccTLD was. They thought that the ccTLD was part of the attack. So it was kind of a shock that they couldn't get pre-action on this. The attack went ahead as planned.

Now afterwards that was a lesson learned that ccTLDs have got to be known to everybody. You can't take it for granted that people know who the ccTLD is. In fact, whenever you try to warn somebody of bad intentions, they have to know ahead of time that you're honestly a good person; you've got to make sure they know who you are. You've got to go through all of this. You've got to go through the network operating groups; you have to make sure there's some organization and you may have to go into some of the more high profile registrars that you have and let them know that you're a registry; you're a good guy; you know what's going on and you have this chance to help them sidestep problems down the road and they work with their operator.

So to kind of wind up the points I have here, DDoS can either be direct to the registry but generally I don't see that as the biggest problem for you. I think most people are doing DNS well enough where you'll see it. The reflection attacks are a concern; they're all over the place; you have to be on the watch for them because they can sneak under the radar. You really don't want to be helping out. You're not really helping them out; you're doing your job and sometimes the more you improve yourself the worse it gets for them. The more you add DNSSEC, the more you add v6, you're giving more ammunition to the reflecting attackers.

And finally, becoming a forecaster – well, you already are able to forecast some of what's out there. You can't predict the attacks, but you can certainly make sure that people out there will know from you, "This is something I've learned," and they'll take it seriously. This may go through the first organizations or the CERTS and so on so that you have to make sure that this community of trust is built around the internet, not just through the internet. And those are the basic steps you can take. Just let people know that you hear something's coming; they trust you for that.

The last thing here, out-of-band ties - law enforcement is important in all of this. In fact, I've heard the comments earlier in other presentations law enforcement does things we don't do; they take care of evidence, they take care of legality, they take care of liabilities. Don't do anything… make sure they know ahead of time everything you plan to do in terms of reaction to bad activity because you could be… just think about this case. What if you're wrong? What if it isn't bad activity? What if it's not illegal? What if it isn't a problem? What if it's somebody else's stuff going forward?

Law enforcement has to be part of the picture in just about everything you do because lawyers matter more than engineers at some point. I hate to say that to engineers but it's kind of true. So thanks for your time. Any questions?

Eberhard Lisse: Maybe I abused the prerogative of the Chair. Last year when I was invited to some prep meeting in Brussels about this, I went to talk to my local law enforcement, to the financial crimes. They had no bloody clue what I was talking about. Absolutely no clue whatsoever, but they realized this; they're very concerned about it because they knew what it meant what I was talking about.

They didn't know about computers. They have two computers in the department.

Difficulty in developing countries is not that easy to get. But they know what fraud is. They know a cook when they see one, so to say. They knew what I was talking about and they were very concerned that they did not have the hardware and the know-how to deal with this if it was happening. Fortunately our prices are so expensive so we're not a target for automated registrations, so it hasn't happened yet as far as botnets is concerned and we have Anycast up to the wazoo so we haven't seen anything happening and it has just absolved.

But even if this doesn't happen you can't close the eyes. You must liaise with local law enforcement. The day will come and it helps especially in smaller countries. Like we remember from the Estonian thing and one of the points I remember very well when they mitigated Digitec it was helpful that the ministers of the permanent secretaries, they all went to the same schools, they're all the same age, they all know each other so communication was easy.

Christopher Davis: And also there was a right meeting going on at the same time and they were external experts doing the Estonia attack. There was a right meeting – actually I was in the country when that happened. I wasn't part of what was going on but…

Eberhard Lisse: No, what I'm saying is I think Hitchcock, Woody, said that's one thing he noticed that these people on the decision making level, they all knew each other very well so that it was easy to get short communications going.

Christopher Davis: To bring Woody in was convenient because he was easily introduced. I should say that my first experience with computer hacking was in the 80s and at the time the police – no idea. There was a book called *The Cuckoo's Egg* and it's a pretty good description of an attack that was in the 80s. I ran into the same exact issue at some point where all the police said, "What's the value of this? What's been lost? If it's less than $1,000, it's a minor crime."

In 2001 I had already been working for DNSSEC for five years and I was asked by some lawyers to have a meeting and it was lawyers from three different countries in the E.U. Cause in 2001 already which was before the spec was last written, they already saw what was happening and they were starting to make plans for what does DNSSEC mean legally. And so sometimes law enforcement – getting them to be aware this is a problem… you don't want to call them up when you want them to arrest somebody. You want them to be involved when they need to learn so that they realize why they would have to at some point arrest somebody.

Eberhard Lisse: My point is police officers like to arrest people. That's what they do and my police officers there – they were quite worried that they did not have the know-how, they did not have the technology, they didn't really understand the finer details but they understood a fraud when they saw one. And they were very concerned that they don't have the know-how and means if and when it comes.

So I full am support this local ccTLD; ccTLDs are small, should talk to their local law enforcement, liaise with them, find out who the chief inspector of the financial crime squad is and know him. Now I've got his number on speed dial and he has got my number on speed dial because I treat his wife – that's a different issue – but the point is you want to be able to talk to somebody who does not think you're a kook. "Wait a minute. He gave us a presentation; he gave us a lecture; he comes every three months and reinforces the lecture; that's the guy that he warned us that it might happen."

It's a good thing to talk to law enforcement that you have sort of a communication established before it happens. Any other questions? I don't want to bore you too much with this. Any other questions? Alright, thank you very much. Let me just think who is the next one. The next one is Morgan Marquis-Boire and Warren Kumari from Google who are going to tell us what happens when you receive the call, "Dude, what happened to my domain?" He's looking for his slides and obviously he keeps them on Google Documents.

Morgan Marquis-Boire: Hi. So this is a short talk about DNS hijackings. My name is Morgan and I work for Google. So I work on the Google Incident Response Team and today I am focusing on compromises external to Google that have nonetheless affected Google's domains. So it's worth noting that none of the compromises that I'm actually discussing have affected Google's domains exclusively, but we're going to be talking about domain hijackings at a level where attackers have thought it prudent or useful to hijack all the really high traffic domains.

So again, standard disclaimer – none of these incidents actually represented a compromise of Google hosts or services and they contain a bunch of real world examples of domain hijackings. So I'm not actually picking on the security of anyone in this room or anyone who I'm mentioning; I'm merely using real examples to highlight the sort of systemic and ongoing nature of a general problem. So again, apologies if you come from any of the countries that I'm going to be discussing. As you guys have probably seen, a bad security day can happen to anyone; it's happened to Google.

So this is just a bit of disambiguation. I just wanted to sort of preface how I'm going to be talking about DNS hijacking. So various people hijack your DNS for kind of pseudo legitimate, I guess even desirable means for them. So hotels do it so they can redirect you to their portals or use their wireless internet services, trying to get you to agree to terms and conditions.

So these types of DNS manglers generally get you to click a box or sign an agreement before they mess with your traffic. What I'm going to discuss today is the people that obviously don't get you to do this before they engage in this type of behavior. There's a bunch of people doing this and they all have kind of… they have differing motivations.

So as I just mentioned, there's people doing it for sort of legitimate or at least legal financial reasons, then there's people doing it for illegitimate financial reasons. Redirection of traffic for the purposes of click fraud or monetizing through advertising has been done using DNS malfeasance for big money.

So the FBI arrested a group of Estonians for profiting off this type of behavior for profiting to the tune of $14 million – this scheme dates back to 2007 and made use of a common botnet trojan to divert web traffic from its intended destination to that of advertisers who paid for traffic delivery and most of the web traffic advertisers thought that this was legitimate traffic.

So the malware at the center of the scam – known as Operation Ghost Click – was a DNS changer trojan. When it was installed, it redirected a host DNS request to a server and basically took control of all the outbound traffic from the affected system. The botnet in DNS servers were controlled by an Estonian company called Rove Digital and its hosting subsidiary ESTHost.

Now amusingly, or perhaps not amusingly, the company even operated its own domain registrar – EST Domains – until it was taken down in 2008 when it lost its ICANN accreditation after the CEO was convicted of credit card fraud.

So DNS hijacking – this is kind of like I guess a sort of more new age form of monetizing, but people have used it for very traditional types of fraud as well. As well as making money *via* traffic redirection, old fashioned credit card harvesting also gets performed using DNS hijacking.

CronoPay was a domain for Russia's largest online payment processor and this was hijacked on Christmas in 2010 and it redirected hundreds of unsuspecting visitors to a fake page that stole customer financial data. CronoPay's domain was transferred to Network Solutions and the DNS servers were changed to anotherbeast.com. The attackers then, *via* this fake page they redirected people to, collected roughly 800 credit card numbers from customers visiting the sites which were used much like Pay Pal to make payments to various Russian businesses. The hackers also stole and posted at least nine of the SSL keys that CronoPay used to sign their SSL certificates.

So moving on from the money, another common motivation for DNS hijacking is censorship. So a lot of stuff has been written about China's DNS hijacking and censorship mechanisms. Basically what they do is they steal your DNS when you go to sites that serve content which they view as possibly undesirable for whatever reason and return content that is more aligned with whatever political agenda that's going on. A lot has been written about this so if you are sort of very interested, I'll leave that as an exercise for the listener.

But to continue along the political vein, hacktivism and defacement are exceedingly popular reasons for DNS hijacking, in fact, sort of recent articles – Wired, the Register and so forth – this year described DNS hijacking as the hactivist's second most popular tool after DDoS.

Some of you guys might remember this. On the 18th of December 2009, this is what the front page of twitter.com looked like. Now in this specific case, their DNS provider had been hacked and their website pointed to a third-party page which had been set up to look like either the Iranian Cyber Army took credit for it or someone taking credit as them. So this is obviously specifically to bring attention to political causes in Iran and so forth.

A sort of compound motivation – which I've listed both – sort of phishing and account access – is where things sort of get really interesting and this is where people harvest user credentials. Now obviously there can be various variations for doing that. However, when this happens at a massive level, the motivations frequently include state surveillance or sort of persecution of people that sort of disagree with certain political agendas.

In this case one of China's largest ISPs, China Net Com, had its DNS servers compromised. Chinese Google users were redirected to a fake page which loaded mail from a malicious domain which was mail.google.com-isifmail-serverlogin.bej900.ndsns01.com which is clearly not a legitimate Google mail domain. The motivation here was credential harvesting of user accounts.

As I mentioned, the consequences for the victims of this can be quite severe, especially when this occurs at a nation's state level. Some of you guys might remember this cause it generated a substantial amount of press. On Christmas 2010, which incidentally was the same time that Chrono Pay was being hijacked as I mentioned, this was a bad time for DNS hijacking – people noticed very large organizations – Google, Facebook, etc. – that someone appeared to be attempting to steal an entire county's worth of user names and passwords.

It looked like .tn had been hijacked and people were harvesting Facebook and Gmail logins using fake login pages. So this is what Google Tunisia looked like at that time. Obviously there's a bunch of things wrong with this page. For a start, Gmail is not a php. If you look in the URL bar, you can see that ServiceLoginAuthservicemai.php was definitely not what Gmail should look like. Additionally, Gmail should not be throwing php errors and run off a Windows box – C:\ProgramFiles\EasyPHP, blah, blah, blah.

So what I'm going to be talking about today specifically are these types of mass domain hijackings and especially what I've seen a lot of is hijacking occurring at a registry level. This is a highly effective way to perform all sorts of compromises for all sorts of reasons. Naturally gaining control of DNS for the DNS records from an entire country gives you a lot of avenues.

Now frequently Warren's meetings interrupt my talks. This is actually the first time this has happened. So frequently when we actually see this type of hacking – a ccTLD registry compromise – what we see is we see the redirection of traffic from most of the high impact, high traffic domains under that country code.

So while the examples that I'm going to be talking about are mostly I'm going to be talking about Google. In most of these cases we saw compromises occurring for the records of Facebook, Bing, Coca-Cola, IBM, Microsoft, etc. etc. were also changed.

So about three years ago we started actually tracking this and so we wrote some really simple code and what it did was it checked the DN records, DNS records, the WHOIS records and all that sort of thing for a variety of high profile Google domains. And the software started firing a lot more often than I would have imagined and I'm not sure if anyone in this room wants to hazard a guess at how often a country gets its ccTLD registry hacked. But basically in the last three years I saw about 16 of these.

So in 2009 Morocco, Tunisia, Tajikistan, Ecuador, Kenya and New Zealand were hacked and then in 2010 Uganda, Puerto Rico, Denmark. 2011 – Suriname, Malawi, Congo, Guadalupe, Fiji and Bangladesh. 2012 we've seen Nepal. So I haven't put this list up here because I'm accusing anyone of being particularly terrible with their security, although there is one person on that list that was hijacked about three… four times actually within the space of a month and that is just bad, actually and I'm not afraid to say that.

But for the rest of them, I pinned this list up here just because I want people to understand that this is a systemic and ongoing problem. It's not up here because everyone who has bad security is up here; it's up here because this is going to keep going. There is no sign that the security of registries has gotten a lot better in a mandated fashion over the last three years so there is nothing to indicate we're not going to see exactly the same type of behavior over the next three years and in fact, with all the new gTLDs, we're probably going to see more of this behavior.

So how does this actually happen? In much the same way as regular compromises happen – software has bugs; people don't know how to sanitize input that they're programs take; long strings appear to be difficult for people to deal with; malformed strings – all that sort of stuff. Even people that run registries reuse their passwords in places that they shouldn't. People get passwords hijacked by logging into kiosks, using their computers in the roach motels of the internet, airports kiosks and that sort of generally bad place.

Social engineering has been a significant problem for several of these high profile hijacks. We have had domains hijacked – very high profile domains – hijacked for countries. People have received emails asking for DNS records to be changed - spoof emails asking for DNS records to be changed and they have done this. People have received requests *via* web forms asking for DNS records to be changed and they have done this. Obviously this is not great.

Bribery and coercion – when as I mentioned, the types of attackers that are interested in this sort of thing are brought to play, this can't really be ruled out and actually the most common technical attack I've seen is SQL injection. For people in the room who aren't intimately familiar with SQL injection, I have put up a humorous comic which describes it which many of you may have seen before.

Basically the problem is that if you have a webpage which receives input, and doesn't sanitize it properly, you may be able to craft input in a way that it is natively understood by the backend server to execute SQL commands. This may be inserting or changing domain records or it may be dropping the table of the student database as per the cartoon.

As I mentioned before though, probably the second most common technique we saw was social engineering. Now I follow a Tumblr called targetedemailattacks.tumbler.com. This is an example of one which is the most unsettled piece of social engineering I think I've ever seen. Generally what you're attempting to do is you're attempting to convince a person to perform an action for you *via* sort of duplicitous means. What the person has done in this email is they've sent someone a thing saying, "Please click on this link." Now I didn't click on that link but I'm pretty sure it's malicious. This is an example of social engineering.

So the effects that we're seeing from these types of compromise – so fortunately the most common thing that I see is web defacement and this is either by hackers or hacktivist groups and they mainly appear to be doing this for bragging rights and I'm gonna show you some of their work in just a second.

Visibility for political causes is a common one for the hacktivist groups and as I mentioned before, monetizing *via* spam, and more scary is mass user credential harvesting.

So this is Google Tunisia. This was hacked by someone called Cyber Mafia Crew. He has [Greeks] and a smiley face and that sort of thing. This one is – I think this was Guadalupe – it' a bit small on the screen. So again, this was actually hacked by Thehacker, Crazy King and so forth. This person is quite keen on taking credit for his hack and again shout-outs to his buddies.

This person compromised Google Bangladesh, redirected the site to a third-party server. Sorry, I should actually clarify that. As I said before, compromised the Bangladesh ccTLD registry and then redirected the Google Bangladesh domain at this website.

This is the same thing for Guadalupe. We actually see a political message here – Gaza in our hearts. This is from a Moroccan hacker group. This one is cd – cd is Congo. So Congo's ccTLD registry got compromised. This is a political message for Tunisia. The DNS for Denmark was compromised and this was redirected to the servers hosted at one.com. So the interesting thing here was one.com is a legitimate third-party hosting provider. They had been compromised and someone set up the bogus name servers there which Google.dk got pointed to. We noticed but by the time we'd been in contact with them, they had noticed by sheer volume of the traffic that had started flying towards them. And, in fact, they even posted this message on Twitter saying, "Oh my God, for some reason, Google.dk is pointing at us. This is not great. We've contacted Google; we're trying to fix this now." And like four people had retweeted it at that point, so obviously it was some sort of news.

And that actually brings me to kind of the public nature of these attacks. While they can be highly effective, both in terms of money theft or credential theft or visibility for your political cause, out of necessity DNS is public and when you redirect the traffic for some of the largest sites in the whole country, people are gonna notice.

So Puerto Rico's DNS got compromised. This article got posted at the time the attackers went after, as I mentioned, all the high profile sites – Google, Microsoft, Yahoo, Coca-Cola and a bunch of other big companies – Pay Pal, Nike, Nokia, etc., etc., etc.

So I kind of want to point out that this is actually an example of the best case scenario press that you'll receive if your ccTLD gets hacked, registry gets hacked. "Hackers temporarily seize control of Google Morocco domain name" shows that the person who wrote this article actually understands something about DNS and the way it works, right? Whereas, frequently what you'll get is something like this, which is just kind of an example of bad reporting. I kind of hope in some ways that reporters will take this to heart as well.

Now this is another incident which occurred recently involving mass credential hijacking and created a significant amount of press but it didn't actually involve DNS hijacking, but it could have. I mention this because at present the problems in the certificate authority in the [SSL] industry are kind of compounded by insecurities in DNS at a country level.

The Diginotar incident which occurred on July 10, 2011 – a wild card certificate got issued for Google by Diginotar. This certificate was used by unknown persons in Iran to conduct a man in the middle attack against Google Services.

On August 28 certificate problems were observed by multiple internet providers and according to a subsequent news release, Diginotar detected an intrusion and admitted their fraudulent certificates had been created for domains of Yahoo, Mozilla, Word Press, the (inaudible) Projects and so on and so forth.

The same hacker that stole certificates from Diginotar also claimed responsibility for stealing certificates from Comodo which is another certificate of authority; StartSSL another CA who admitted to having key material stolen and recently the CA Trustwave admitted to mincing a certificate specifically to allow the man in the middle sort of thing of HGPS traffic.

I mention this because part of the entire structure of HGPS is the sale of TLS and Global Certificate Authorities is that DNS is being honest with you. DNS is part of Trusted System. So obviously you have real problems if someone or something has modified your host file or poisoned the DNS. Sadly much of the trusted communication in both the CA and the DNS worlds occurs over plain text email. The type of DNS hijacking I've described can not only be leveraged to mass harvest credentials but also fraudulently obtained certificates and domains, especially if people start off hijacking mx records that forge emails between parties that provide these types of services.

So why should you or ICANN care? Well, having a ccTLD registry hacked is really embarrassing. It's really terrible press and as I said, press really tends to notice when all the most high profile sites in the country all of a sudden point at another page. And as we've discussed the motives of these groups range from basic fraud to mass surveillance in regimes with less than perfect human rights records. In these types of places things can end very poorly for people that have their credentials and accounts hijacked. The trust in the DNS system is an integral part of the internet and the trust in registrars and ccTLD registries is inherited from ICANN.

So what am I actually suggesting that people do? I'm suggesting that we actually have mandated regular security audits for registries. I think that a minimum base line of security needs to be mandated and described in the registrar program.

I also have seen several types of registry software available. I think mandating a registry in a software that is actually heavily audited by the security community and therefore isn't plagued by SQL injection and other types of attack is a really good idea and obviously, DNSSEC, the sooner the better.

So any questions? Yes? The guy from Emerging Threats?

Male: Hi. Your last point about registering a box I think is very interesting. If you look at some of the registrars like [Direct I], for example, when you're talking about EST domains which you know, I inherit all of that, right? They're got logic boxes, I think it's called, which is a registrar reseller type of deal box. I think that having some sort of common code is a really good idea. I just don't know… do you have ideas on how you would start that? Would you make it an Open Source project? Would you…

Morgan Marquis-Boire: Yeah, so I've discussed with various people the way you actually get the security community to audit your code is either you pay them and you pay people lots of money every time you make a code revision and you pay them hundreds of dollars an hour and you make sure they're good. Or you have code that is used ubiquitously for really important parts of the internet.

For instance, Apache is incredibly heavily audited code; Bind is now heavily audited code. This is because they're important and a lot of security eyes have been over this Open Source code a lot of times. If you have a similar type of thing for DNS registries, then I think that would be a good start.

Eberhard Lisse: You haven't been working with many ccTLD registries about it. There is no way of mandating this because if I can trust (inaudible) – they cannot mandate anything. It's not an excuse. Many of us run corporate (inaudible), seriously audited and we haven't been aware of this. Some run FRED; some are very responsible; some are less responsible. I looked but the name is not on there. The point is to say we call this [blue-eyed] in Germany, to say they must be audited, they must be mandated. I can mandate whatever I want; nobody cares because we have no legal relationship to ICANN. There is no country; there's no registrar agreement between the ccTLD… I pre-date… .na predates ICANN by 10 years. I don't care much mandate, but some upstarts want to mandate me.

Morgan Marquis-Boire: Thank you for calling me an upstart. I don't think mandate… obviously not mandating for ccTLDs – you know, recommending for ccTLDs or the new gTLDs – that could be a place where there could be and if that is mandated and the code is well audited there, then the ccs can benefit from that – free, audited code if you want to use it. So that was potentially worded poorly.

Male: I've got a question. I really like the idea of trying to get some sort of audited code and have some deeper questions. You know there are various kinds of audits and you know with the DNSSEC stuff, I'm a component of paying the big money and making things audited. As you pointed out, most people are not going to be able to afford that and are not going to do that and there's no way we can mandate that.

But then you said that things like Bind were heavily audited and Apache and I know that some versions of Linux happen to be pretty carefully audited. Can you tell me more about that because I'm intrigued at this idea and I'm certainly not stepping up to the plate because it's above my pay grade, but it sounds like the kind of thing ICANN might be interested in helping on it.

Morgan Marquis-Boire: So I can't unfortunately speak for the entire security community and if I say if you publish it they will audit it. Having said that, what I mean is that the types of code that tend to attract the security community to audit it create feedback (inaudible) tend to be as I said heavily used and important code. Getting somebody to audit your code which three people use is never gonna happen.

But for instance, people also offer bug bounties these days. For instance, we have… Chrome has a vulnerability bounty program where – as opposed to paying the big, big monies, the prize that is commonly paid out is $1,337 which is sort of a token recognition of… And we publish their name and say thank you very much; this was a critical vulnerability discovered by this person. And so these types of community initiatives have proven successful in attracting community auditing and that sort of thing.

Male: Sounds cool. I just have one other question then I'll be done. You spoke earlier about the CronoPay? How did they get the SSL private key?

Morgan Marquis-Boire: Oh, so I don't actually have any knowledge other than public knowledge of this security incident. So it's poss… I mean, given that they actually own CronoPay's DNS, they could simply have hijacked the mail of… you have the horizontal and the vertical once you've got it so…

Male: I guess I'll steal the mic. So actually one other thing that I guess I should have mentioned is there have been a number of ccTLDs that we've had friends with [hat check] deals like Steven and Nigel and a few other, who, we've asked them nicely to please watch for this sort of stuff and a couple of times people have sent us stuff like, "I got this really weird email asking me to please transfer your name server somewhere else. I said no, it wouldn't have worked even if I hadn't been doing it." And it's often just interesting to know. That's always nice when people are willing to do that sort of thing.

Eberhard Lisse: Have any of these hijacked cc or tlds used CoCCA or FRED, one of the Open Source registries, do you know?

Morgan Marquis-Boire: I don't know. I mean it said there were 16 so I actually don't know what software each individual…

Male: I believe so. I'm trying to remember at least… I believe at least one of them was running FRED… I think it was FRED, not CoCCA. What the actual issue was is they had a web-based front end for WHOIS and that had a SQL injection attack. And then once the people had that, they then managed to leverage that for additional stuff. So while in general the registry systems themselves might be okay, often in a fit of enthusiasm people put up additional stuff in front of that and they do it in a rush and don't necessarily sanitize the input.

Eberhard Lisse: I'm just interested whether I'm one; we run the (inaudible) but we don't do that kind of thing – not that I understand them but probably that's the reason why we don't do that. Anyway, thank you very much. The next one is Francisco Arias who is going to talk to us about some new next generation WHOIS.

Francisco Arias: Thank you. Hello everyone. I am Francisco Arias. I work with ICANN. I am on the technical side of ICANN. I used to be on the other side where you are. I used to work for .mx. It's good to be here back and see some familiar faces.

I'm going to talk about replacement of WHOIS. This is a topic that some of you may have heard that isn't ready. Something going on in (inaudible) releasing the idea. So why we need to replace WHOIS – I guess most of you already will have some ideas why we need to replace it. Some of the main reasons why we think it should be replaced is, for example, the lack of international registration support, the lack of having a way to (inaudible) the code that is being used; the lack of standardization in the query response and error message from the protocol and also some registries would like to have authentication and access control on the information. Some details use for example, .name, .tell, to the WHOIS information.

So this is something that is difficult to do with the WHOIS protocol. Actually what they are doing is they only offer that on the web-based interface. Another thing that is also missing is mechanism to discover the authoritative servers for the WHOIS upon a specific (inaudible).

So what's the options we have? Well, we had a previous option that was specified in the IDF - you may have heard of it. However, it seems nobody's interested in implementing. Many people say it's too complicated. So there is one option that has emerged and it's already being used by the RIR community, ARIN, RIPE already have a web-based interface, a RESTful WHOIS interface for the database and some other RIRs are already working on this.

So what are the benefits of using RESTful WHOIS? Well, this is something that many people already know how to do; it's simple. There are a lot of libraries for client on server side. You can use their web browser if you have XML with CSS, you can easily have it for the human family interface let's say. So you could have both standardized output and you can have also the human family output in one service instead of having to offer two or three and the web-base.

On the sign side you would be using these kind of properties that may be interesting to have. I won't go into detail into this since we have not much time. In ICANN a colleague and I did a short pilot on how it would look RESTful WHOIS for a domain name registry. This is the [URI server] that we were thinking of using… sorry, the one that we use. As you can see, it's very simple; there is nothing really interesting to talk about here. The address for the domain name, you put the name and you could put it in the A-label, the U-label.

For the context we were using the ID. So however the areas we're talking with them and they were using a slightly different interface and we tried to accommodate that see if that makes sense in the domain name side. This is the kind of interface they are using or shall I say the interpretation that they did of that interface which is basically the same, just you have a predicate after the name.

So you can ask for example for the context of the name of the registrar or for a specific contact with this predicate. And similarly you have that for the context and the host. So this is I think… the other thing is that after four months there is still discussion in the group what should be the output format – XML, JSON or something else.

So I mentioned there is already a mailing list in the IETF which some of you already participate in. Basically the message here is if you haven't seen this many lists, you probably want to take a look and see if there is something of interest to you.

What is the main idea here is to use the setup requirements that were already defined for the ID support and the increased requirements, to use that as a base, but now to build RESTful interface for the WHOIS data.

We already held a BOF in Taipei, the IETF in Taipei last year. There was a strong presence of the RIRs there so it seems like the numbers side of the WHOIS let's say, there is enough support to have our working group with them in scope. The discussion is what will be the domain name side in scope (inaudible). A number of registries from the gTLD and the ccTLD side have shown support in the mailing list but it seems like the AV, the other actor, would like to see more support on the domain name side. So if this is of interest to you, please go to the mailing list and participate.

So, yes, there will be another BOF in Paris. The IETF part is in a couple weeks and we will try to reach agreement on the charter for this working group. Basically this will be the last call for having this working group in the IETF. We have been told it's very unlikely to have a [tier work], so basically we need to make it happen in Paris if we want this to happen.

This is the link to the mailing list. I forgot what "WEIRDS" means – worldwide something – I guess it doesn't matter. This is the link - if you are interested, please go there and subscribe and participate in the mailing list. Thank you.

Eberhard Lisse: Okay thank you very much. Some say that one shouldn't do technical stuff on ICANN meeting. In particular [Sabina Doto] who isn't here – I had a little fight with her yesterday about it, but I think today shows – and Andre will probably delve a little bit more into this – that there is a place of doing a technical shop with ICANN. I think even if ICANN has nothing to do with IETF or writing the protocol it's good that we are aware of things like this will be coming. Our software will have to be adapted; we have to write interfaces for the databases, for the back ends and so on. So it's quite a good thing to be kept abreast. In any case, without further ado, Andre will close this up. Questions? Sorry, Ed Lewis has…

Edward Lewis: Ed Lewis, NewStar. On slide 12, you said it was broadly aligned. Can you say quickly how it's not aligned? Is there a significant difference? Is like the beep thing… whatever?

Francisco Arias: No, the only thing I meant is that the idea is to use that as a base…

Edward Lewis: What did you not like about the [crisp] requirements?

Francisco Arias: It's not that. All I'm saying is the IDN Working Group and the draft charter – I don't know if you have seen it – mentions these (inaudible) to be used as a surbase for the requirements. It's not it's something we don't like. We start with that but we are not close to new things or changes. We will leave if there is something we don't like.

Warren Kumari: Warren Kumari, Google. So this isn't really a question; it's more of a statement or a soapboxy thing. In the WEIRDS group, there are already a number of the number organizations or a number of registries participating and they have a draft and a lot of discussion about it and a lot of support. And the sort of naming side of it is largely being treated as second class citizens.

There's been discussions that the naming folk aren't showing up and aren't participating and we don't really see a desire from the naming people to participate. And so we're sort of having a hard time being included in the charter. So if this is stuff that's interesting to you or stuff you care about, please just show up on the list and be like, "I'm from a name registry. I'm interested; I'm willing to review documents. I'll come along and help." And that will actually stop us being the bastard stepchild. That's all.

Eberhard Lisse: Alright, any other questions? So then Andre can close this up.

Andre Filip: Thank you very much. It was a great meeting as usual and we have basically two very interesting topics. First of all is the registry systems and we learned a little bit more about FRED and CoCCA. Now we had the three presentations about FRED and that's probably caused by the fact that we are in a country that uses FRED as the main registry system.

And it was really interesting topic. It's something we need to think about for Prague because Prague is the hometown of FRED and we were discussing internally with Eberhard and others that maybe we could extend the format of the Tech Day to some sort of work shop, but let's see what the Technical Working Group will decide and how it is going to work. We need to think about that and prepare something that will be interesting for you.

The second part was again very interesting. It was about security behind DNSSEC so that's something I was really looking for and I hope you like it as well. So we learned a lot about the DDoS and things like that. And also I saw one sub-topic. It was an Estonian and Turkish presentation which was the registrar identification – something we have never touched on, so that's also a topic that was really not widely discussed and we can save that to pick for some future meetings.

So again, what I have learned from my group at United Nations. Every international meeting is either a success or an outstanding success. I hope you will agree this is an outstanding success and you will help to thank the Chairman for the great job and the presenters for the perfect presentation. Thank you very much. Enjoy the rest of the day.

[End of Transcript]